

Lifted Auto-Context Forests for Brain Tumour Segmentation

Loic Le Folgoc, Aditya V. Nori, Siddharth Ancha, and Antonio Criminisi

Microsoft Research Cambridge, UK.

Abstract. We revisit Auto-Context Forests for brain tumour segmentation in multi-channel magnetic resonance images, where semantic context is progressively built and refined via successive layers of Decision Forests (DFs). Specifically, we make the following contributions: 1) *improved generalization* via an efficient node-splitting criterion based on hold-out estimates, 2) *increased compactness* at a tree-level, thereby yielding shallow discriminative ensembles trained orders of magnitude faster, and 3) *guided semantic bagging* that exposes latent data-space semantics captured by forest pathways. The proposed framework is practical: the per-layer training is fast, modular and robust. It was a top performer in the MICCAI 2016 BRATS (Brain Tumour Segmentation) challenge, and this paper aims to discuss and provide details about the challenge entry.

1 Introduction

The past few years have witnessed a vast body of machine learning (ML) techniques for the automatic segmentation of medical images. Decision forests (DFs) [21,18], and more recently, deep neural networks [11] have yielded state-of-the-art results at MICCAI BRATS (Brain Tumour Segmentation) challenges.

In this paper, we describe an approach that builds upon the framework of DFs, departing from the usual hand-crafting of powerful features based on *e.g.* texture, elastic registration, supervoxels [17,6] with a complementary scheme that is entirely generic and free of additional computations. More specifically, we introduce an efficient node-splitting criterion based on cross-validation estimates that improves the feature *selection* during the training stage. Consequently, learnt features are more discriminative and generalize better to unseen data: we refer to this process as *lifting* (see Section 4). Furthermore, the proposed cost function induces a natural stopping condition to grow or prune decision trees, resulting in an *Occam's razor*-like, principled trade-off between tree complexity and training accuracy. We show that lifted DFs can outperform standard DFs using compact, shallow tree architectures (several dozens or hundreds of nodes).

We exploit the resulting computational gains to revisit *auto-context* [14,15,13] segmentation forests. In particular, we extend the approach with a meta-architecture of cascaded DFs that naturally intertwine high-level *semantic* reasoning with *intensity*-based low-level reasoning (described in Section 5). The framework aims to be practical: the per-layer training is simple, modular and robust. Furthermore, this sequential design enables the decomposition of complex segmentation tasks into series of simpler subtasks that, for instance, exploit the hierarchical structure between labels (*e.g.*, whole tumour,

tumour core, enhancing tumour parts). Beyond auto-context, another contribution is a clustering mechanism (see Section 5.2) that exposes the latent data-space semantics encoded within DF pathways. The learnt semantics are exploited to automatically guide classification in subsequent layers. Finally, in Section 7, we discuss results and details of our BRATS challenge entry.

2 Background: Random Forests for Image Segmentation

Let $I = \{I_j\}_{j=1..J}$ be the set of input channels in a multichannel image, x be a voxel, and $c \in \mathcal{C} = \{1 \dots K\}$ be the class to predict for x . We define $\mathbf{x} = \{x, I\} \in \mathcal{X}$ as the feature vector of x . DF classifiers predict the probability $p(c|\mathbf{x})$ that a voxel x belongs to class c given its feature representation \mathbf{x} . This is done by aggregating predictions of an ensemble of \mathcal{T} decision trees. Simple averaging is typically used, so that:

$$p(c|\mathbf{x}) = 1/|\mathcal{T}| \cdot \sum_{t=1}^{|\mathcal{T}|} p_t(c|\mathbf{x}),$$

where $p_t(c|\mathbf{x})$ denotes the prediction by tree $t \in \mathcal{T}$. The class with maximum probability is then returned (maximum a posteriori probability) as the prediction. Tree predictions are obtained as follows: a feature vector \mathbf{x} is routed along a path in the tree from the root node by evaluating it at every internal (*split*) node n w.r.t. a routing function $h_n(\mathbf{x}) \triangleq [f(\mathbf{x}, \boldsymbol{\theta}_n) \leq \tau_n] \in \{0, 1\}$, and taking the left child n_L if $h_n(\mathbf{x}) = 1$, and the right child n_R otherwise, until a leaf node $n(\mathbf{x})$ is reached. Here $f : \mathcal{X} \times \Theta \mapsto \mathbb{R}$ is a weak learner parameterized by a node-specific feature $\boldsymbol{\theta}_n \in \Theta$, and $\tau_n \in \mathbb{R}$ a node-specific threshold. Examples of weak learners are given in section 3. Each terminal (*leaf*) node $n \in \mathcal{L}$ is paired with a class predictor $p_n(c|\mathbf{x}) \triangleq p_n(c)$ such that $0 \leq p_n(c) \leq 1$, $\sum_{c=1}^K p_n(c) = 1$. The tree then predicts class c with probability $p_{n(\mathbf{x})}(c)$.

Decision trees are trained in a supervised manner from training data $D = \{\mathbf{x}_i, c_i\}_{i=1}^N$ with known class labels c_i , greedily and recursively, starting from a single root node. Training a node n entails finding the optimal feature $\boldsymbol{\theta}_n^*$ and threshold τ_n^* such that the node training data $D_n = \{\mathbf{x}_i, c_i\}_{i \in \mathcal{I}_n}$ is split between left and right children n_L and n_R in a way that maximizes class purity. Specifically, $\psi^* = (\boldsymbol{\theta}_n^*, \tau_n^*)$ and the resulting split $D_n = D_{n_L}(\psi) \amalg D_{n_R}(\psi)$ maximizes the Information Gain $G(\psi; D_n)$, which is defined as follows:

$$G(\psi; D_n) \triangleq \sum_{c \in \{L, R\}} \frac{|D_{n_c}(\psi)|}{|D_n|} \sum_{c \in \mathcal{C}} p_{n_c}(c; \psi) \log p_{n_c}(c; \psi) - \sum_{c \in \mathcal{C}} p_n^*(c) \log p_n^*(c), \quad (1)$$

where $p_{n_c}(c; \psi) \triangleq |\{i \in \mathcal{I}_{n_c}(\psi) | c_i = c\}| / |D_{n_c}(\psi)|$ is the empirical class distribution in the training data D_{n_c} for the child node n_c . The optimum $\psi^* \triangleq \arg \max_{\psi} G(\psi; D_n)$ is found by exhaustive search after proper quantization of thresholds τ_n . Trees are grown up to a predefined maximum depth, or until the number of training examples reaching a node is below a given threshold.

Last but not least, random forests introduce randomization in the training of each tree via feature and data *bagging*. For the t -th tree and at node $n \in \mathcal{V}_t$, only a random

subset $\Theta' \subsetneq \Theta$ of candidate features is considered for training [1,7]. Similarly, only a random subset $D'_n \subsetneq D_n$ of training examples sampled with(out) replacement is used [3]. Data bagging is implemented both at an image level (random image subsets) and at a voxel level (random voxel subsets).

It is important to note that we do *not* make use of class rebalancing schemes. Training samples are often weighted according to the relative frequency of their class. This strategy aims to correct classifier bias in favor of the more frequent class, in cases where there is a large class imbalance. Of course, class rebalancing induces the opposite bias against more prevalent classes, which cannot be avoided in a multilabel classification setting. Section 5.2 discusses an alternative strategy to naturally correct for distribution imbalance.

3 Fast scale-space context-sensitive features

We revisit scale-space representations to craft fast, expressive, compactly parameterized features, as a simple alternative to the popular integral or Haar-like features [19].

Background: integral features. Integral features are based on intensity averages within anisotropic cuboids offset from the point of interest [5]. Cuboid averages are computed in constant time by probing the value of a precomputed integral map at the cuboid vertices [19]. For instance, $f(\mathbf{x}, \boldsymbol{\theta}) \triangleq \sum_{x' \in \mathcal{C}_2} I_{j_2}(x') - \sum_{x' \in \mathcal{C}_1} I_{j_1}(x')$ computes the difference of responses in cuboids \mathcal{C}_1 and \mathcal{C}_2 of size $\mathbf{s}_1 = (s_1^x, s_1^y, s_1^z)$ and $\mathbf{s}_2 = (s_2^x, s_2^y, s_2^z)$, centered at offset locations $x + \mathbf{o}_1$ and $x + \mathbf{o}_2$, in distinct channels I_{j_1} and I_{j_2} . Here $\boldsymbol{\theta} = (j_1, j_2, \mathbf{o}_1, \mathbf{o}_2, \mathbf{s}_1, \mathbf{s}_2)$ is a 14-dimensional feature.

Proposed scale-space representation. During node training, it is crucial for sufficiently strong weak-learners to be computable within the budget allocated to feature sampling and optimization. Therefore, reducing the feature parameterization while maintaining expressiveness is key. For integral features, the sophistication of probing anisotropic cuboids with a continuous range of edge lengths comes at a cost w.r.t. parametric complexity. We restrict ourselves to a small *finite* range of *isotropic* averages. We augment the original set of c input channels with their smoothed counterparts under separable Gaussian filtering at scales σ_1, σ_2 etc. Given s scales, $f(\mathbf{x}, \boldsymbol{\theta}) \triangleq I_{j_2}(x + \mathbf{o}_2) - I_{j_1}(x + \mathbf{o}_1)$ computes the difference of responses in different channels at different scales and offsets from the voxel of interest ($j \in \{1 \dots c \times s\}$ flatly indexes channels and scales). Here $\boldsymbol{\theta} = (j_1, j_2, \mathbf{o}_1, \mathbf{o}_2)$ is an 8-dimensional feature.

A single point is probed for every 8 cuboid vertices probed under integral features, as well as circumventing many boundary checks. For all practical purposes ($s = 2, 3$), `byte[]` storage of scale-space maps limits the memory overhead relative to integral maps (`short[]` storage).

Fast rotation invariant features. We can go beyond *directional* context and account for natural local invariances with fast, multiscale, approximately rotation invariant feature. Let $\phi_1 \dots \phi_{12}$ stand for the coordinates of an axis-aligned, centered icosahedron of radius r . Denoting by $\boldsymbol{\theta} = (j_1, j_2, r)$ the 3-dimensional feature, $f(\mathbf{x}, \boldsymbol{\theta}) \triangleq \text{Median}_{v=1}^{12} |I_{j_2}(x +$

$|\phi_v) - I_{j_1}(x)|$ gives a robust summary of intensity variations around point x and probes 13 points only.

4 Lifting Decision Forests by Minimizing Cross-Validation Error Estimates

4.1 A cautionary look at Information Gain maximization

We follow the notation introduced in Section 2, but drop the node index n for convenience. Information gain maximization w.r.t. (feature, threshold) parameters $\psi = (\theta, \tau)$ can be shown to be equivalent to a joint maximum likelihood estimation (MLE) of $\phi \triangleq (\psi, p_L, p_R)$, the node parameters and children’s class predictors. We omit the details for the sake of brevity. Essentially, decision trees are usually grown by greedily, recursively splitting leaf nodes by likelihood maximization:

$$\phi^* = \arg \max_{\phi} \mathcal{C}(D; \phi) \triangleq \arg \max_{\phi} \frac{p(D; \phi)}{p^*(D)}. \quad (2)$$

In Eq. (2), the denominator is the data likelihood using the current leaf node predictor, whereas the numerator is the data likelihood when splitting this node with parameters ϕ into left and right children. The denominator is constant w.r.t. ϕ , as optimization of the current node has precedence in the recursive schedule.

MLE runs the risk of overfitting. At a node-level, weak learners with poor generalization may be selected. The deeper the trees, the more likely it is to happen, since the training data is split between an exponentially increasing number of nodes. At a tree-level, the lack of principled method for controlling model capacity negatively impacts generalization. Indeed, the information gain is strictly positive (the likelihood ratio of Eq. (2) is > 1) as long as: 1) training samples remain at the node of interest, and 2) the data distribution is not pure. As a result, trees generally grow to the maximum allowed depth, with little control over generalization. Medical image segmentation tasks often require large trees of weak learners to be grown (tree depth 20–30, millions of nodes). Due to computational constraints, few such trees can be grown (a few dozens at most). For this reason, model averaging across randomized trees is insufficient to balance tree overfitting. As an efficient alternative to MLE that can be directly used to control tree growth and generalization, we propose to maximize the predictive score as obtained from *cross-validation* estimates.

4.2 Maximizing Cross Validation Estimates of Generalization

We derive Cross-Validation Estimates (CVE) of the predictive score as follows. At any given node, the (potentially bagged) training data $D = D^V \amalg D^T$ is randomly divided into two disjoint subsets, a tuning subset D^T and a validation subset D^V . The optimization problem is then defined as:

$$\begin{aligned} \phi^* &= \arg \max_{\phi} \frac{p(D^V; \theta)}{p^*(D^V)}, \quad \text{s.t.} \\ (\tau_{\theta}, p_{\epsilon}(\cdot; \theta)) &= \arg \max_{\tau, p_{\epsilon}} p_{\epsilon}(D_{\epsilon}^T; \phi) \end{aligned} \quad (3)$$

where $p(D^V; \theta) \triangleq p(D^V; \theta, \tau_\theta, p_\epsilon(\cdot; \theta))$. The key change is that parameters are now constrained to be tuned on D^T , whereas the final feature score is computed on D^V . While a k -fold estimate could be used instead in Eq. (3), the *hold-out* procedure has the benefit of efficiency and added randomness.

Key to the proposed approach is that the quantity in Eq. (3) takes values ≤ 1 whenever no candidate split yields superior generalization to the current leaf node model. Based on this, we implement a greedy scheme to control the tree complexity, where branches that do not increase the score are pruned in a single bottom-up pass as post-processing, similarly to [12]. We further use a simple heuristic to drastically reduce training time, growing trees *in stages* up to any desired maximum depth, successively pruning score-decreasing branches and regrowing remaining, non-pruned ones.

5 Auto-context forests for brain tumour segmentation

We investigate cascaded DF architectures, made of layers of DFs partially or fully connected via their output posterior maps. This architecture naturally interleaves high-level *semantic* reasoning with *intensity*-based low-level reasoning. We demonstrate this via two ideas: a) *auto-context*: allowing downstream layers to reason about semantics captured in upstream layers, and b) *decision pathway clustering*: latent data-space semantics are revealed by clustering decision pathways and cluster-specific DFs are trained.

5.1 Building and training Auto-Context Forests

The process of cascading DFs is illustrated in Fig. 1. Since DFs rely on generic context-sensitive features that disregard the exact nature of input channels (cf. Section 3), we simply proceed by augmenting the set of input channels for subsequent layers with output posterior maps from previous layers.

Layers are trained sequentially, one at a time in a greedy manner (following Sections 2 and 4.2). Specifically, for the BRATS challenge, class labels follow a *nested* structure: the whole tumour (WT) consists of the edema (ED) and tumour core (TC). The tumour core itself is subdivided into enhancing tumour parts (ET), and other parts of the core that are only indirectly relevant to the task: the necrotic core (NC) and non-enhancing remaining parts (NE). Usually, these labels would be interpreted as mutually exclusive classes (ED, ET, NC, NE and the background BG) so as to formulate the task as a multilabel classification problem. Instead, our proposed framework directly uses these hierarchical relations. While many variants can reasonably be built, the final architecture that we used for the BRATS challenge consists of layers of binary DFs, alternating between predictions of WT, TC and ET (section 6.2).

5.2 Exposing Latent Semantics in Decision Forests for guided bagging

Given Auto-Context Forests, a natural idea is to progressively refine the region of interest (ROI) after each layer, starting from an initial over-approximation of the ROI (*e.g.*, the full image). Downstream DFs are trained on more refined ROIs that exclude irrelevant background clutter, thereby increasing accuracy. Such an approach runs the risk of

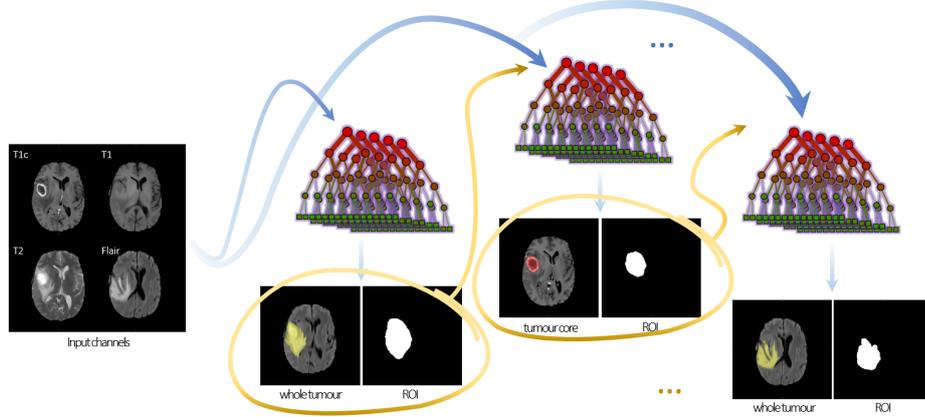


Fig. 1. Auto-Context Segmentation Forests. In this schematic example, layer 2 solves a segmentation task distinct from that of layers 1, 3 but the interleaving allows to exploit joint dependencies.

excluding false negatives, creating a trade-off between coarser ROIs (high recall) and tighter ROIs (high precision). We investigate a complementary strategy to circumvent this limitation. ROI refinement remains as a computational convenience.

The proposed approach exposes and exploits the latent semantics already captured within a given DF as follows. Each data point is identified with the collection of tree paths that it traverses. A metric d_{DP} is defined over such collections of tree paths (*decision pathways*), assigning smaller distance between points following similar paths across many trees, and data clusters are identified by k -means w.r.t. d_{DP} . Then, cluster-specific DFs are trained over the corresponding training data. At test time, data points are assigned to the cluster with closest centroid and the corresponding DF is used for prediction.

Our key insight is that data points that are clustered together will share common underlying semantics, as they jointly satisfy many predicates (see Fig. 2). A wide range of metrics can be designed and for the sake of simplicity, we define (given a collection \mathcal{T} of trees):

$$d_{DP}^{\mathcal{T}}(\mathbf{x}_i, \mathbf{x}_j) \triangleq \sum_{t=1}^{|\mathcal{T}|} \left(\frac{1}{2}\right)^{\text{depth}_t^{\mathcal{T}}(\mathbf{x}_i, \mathbf{x}_j)}, \quad (4)$$

where $\mathbf{x}_i, \mathbf{x}_j$ are two points, and $\text{depth}_t^{\mathcal{T}}(\mathbf{x}_i, \mathbf{x}_j)$ is the depth of the deepest common node in both paths for the t -th tree ($+\infty$ if the paths are identical).

6 BRATS challenge: framework details

6.1 Training dataset (BRATS 2015)

The BRATS 2015 dataset was available to participants of the challenge. It contains 274 images together with their ground truth annotations. One of the interesting aspects of

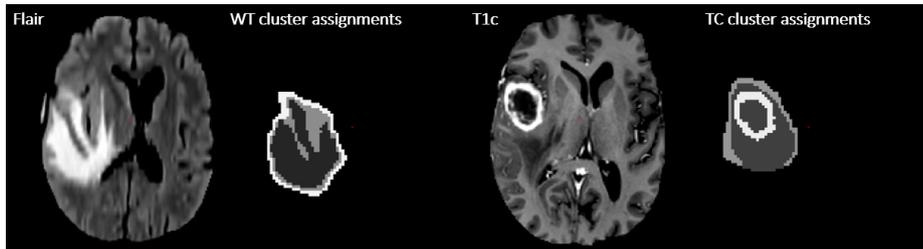


Fig. 2. Voxel cluster assignments for an example subject. Clusters for WT and TC binary forests are learned independently. 4 clusters are used for each and assignments are colour coded in gray levels. On this task clusters naturally appear to relate to "boundary" regions of higher uncertainty and higher certainty regions ("inside" the tumour or the BG).

the dataset is the nature of ground truth annotations: 30 images (from the BRATS 2013 dataset) were manually annotated, and the remaining were annotated using a consensus of segmentation algorithms [9]. While the ground truth is often of good quality, we note with interest that the consensus of algorithms generally fails at correctly labelling post-resection cavities as in Fig. 3 (bottom row). This is likely due to the fact that there is only *one* such training example in the original BRATS 2013 dataset (Fig. 3, top row).

We paid particular attention to such training examples. For these cases, we favoured a qualitative, visual assessment of correctness over quantitative metrics (DICE overlap or Hausdorff distance) when tuning our pipeline. These cases and similar observations motivate the two following choices: 1) An unsupervised SMM/MRF (see below) is trained on the 30 manually annotated BRATS 2013 images, to initialize the segmentation pipeline. For the background class, SMM weights are spatially varying, so that the model proves reasonably effective to disambiguate potential post-resection cavities from, say, ventricles, and 2) 70 images from the BRATS 2015 dataset ground truth with high quality annotations are chosen and used for training of the final model. While leading to a slight decrease in quantitative performance of the algorithm, it also qualitatively somewhat improves segmentation results (Fig. 3, last column). The same qualitative observations are made on the BRATS 2016 test set.

6.2 Pipeline, model and parameter settings

Preprocessing. Image masks are defined from the FLAIR modality, masking out 0-intensity voxels. The intensity range is standardized: the distribution of voxel intensities within the mask is normalized to a common median and mean absolute deviation by affine remapping. As a mostly implementation specific step, we further window intensity values to make threshold quantization easier when training DFs: intensities are thresholded and brought within a byte range.

Initialization: SMM/MRF. A Student Mixture Model (SMM) with Markov Random Field (MRF) spatial prior is used to locate the region of interest (ROI) for the whole tumor. The likelihood for each of the five mutually exclusive ground truth classes is modelled by an SMM with spatially varying (BG) or fixed (other classes) proportions [2],

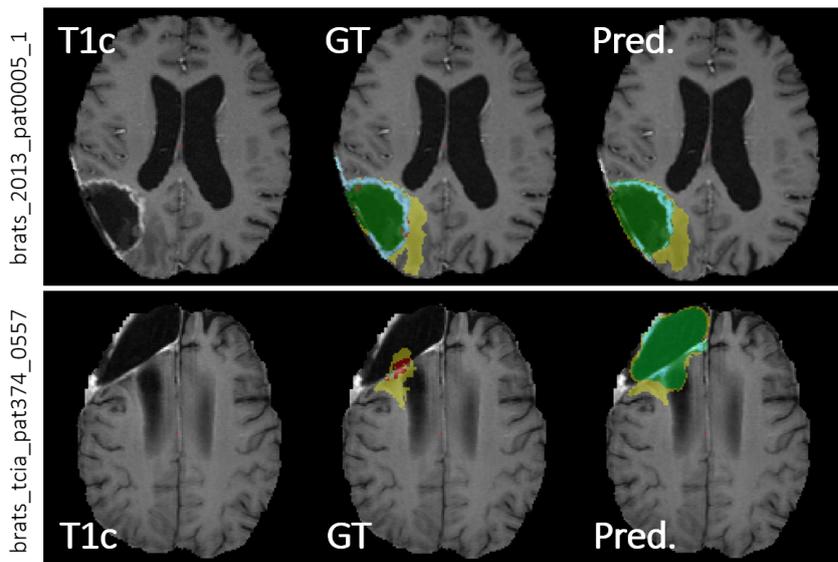


Fig. 3. Example images with resection cavities in the BRATS 2015 training set. The first column displays the gadolinium enhanced images, while the second and third respectively display the ground truth annotations and algorithm predictions. Top row: manual annotation. Bottom row: consensus annotation.

as a suitably modified variant of [4]. An MRF prior is assumed over BG, ED and TC. The model is similar in spirit to [20,10]. It is unsupervised: the current implementation does not use white/grey matter and cerebro-spinal fluid labels, but the resulting components for the background SMM are highly correlated to those labels. Variational Bayesian inference is used at training and test time. The MRF defines fully connected cliques over the image, with Gaussian decay of pairwise potentials w.r.t. the distance of voxel centers. For this choice of potentials, the dependencies induced by MRF priors in variational updates can be efficiently computed via Gaussian filtering. Inference over $3D$ volumes is fast at training (seconds or minutes) and test time (seconds).

Auto-context architecture. 9 layers of binary DFs are cascaded, cycling between WT, TC and ET probabilities. All layers use the original, raw image channels. The input of the first layers is augmented with probability maps from the upstream SMM/MRF, the subsequent layers use probability maps output by previous layers. For instance, the second TC layer uses the output of the first TC, ET layers and of the second WT layer. In addition, the prior probability “atlas” maps returned by the spatially-varying background SMM/MRF model are passed to the first three layers. Many variants of this architecture were informally tested without a significant effect on accuracy.

ROI refinement. For computational convenience, subsequent layers are run on ROIs rather than the full image. For instance, the second BG vs. WT binary forest only tests

points within the mask provided by the first BG vs. WT forest. Similarly at training time, the second layer is trained on subsets of image voxels within the respective image ROIs output by the first layer. Each layer uses masks obtained as dilated versions of the segmentation masks output by the previous layer (dilated resp. by 15mm, 10mm, 5mm).

Parameter settings. DFs come with a number of parameters, many of which do not seem to strongly affect the pipeline accuracy after informal experiments. Five feature types are used: intensity in a given channel (respectively, at scale s), difference of intensities between the voxel of interest (VOI) at scale s_1 and an offset voxel at scale $s_2 > s_1$ in a given channel, median of the intensity difference (respectively, absolute difference) between the VOI at scale s_1 and the radius- r icosahedron vertices (scale $s_2 > s_1$) in a given channel. Between 100 and 200 candidate features are sampled per node. We use 2 scales: 1mm and 2mm for the first layers, 0.5mm and 1mm for the second and final layers. The offset along each direction (or the icosahedron radius) is sampled uniformly between 0mm and 50mm. The range of feature responses is quantized using 50 thresholds. The maximum tree depth is set at 12, and is seldom reached. The number of decision pathways clusters (section 5.2) is set to 4. They are created using the first layer of WT, TC, ET (separately for each classification task). 50 trees are trained per layer per cluster-specific DF¹. The subsampling rate for data bagging is adjusted based on the desired computation time. At each node, the training voxels from 25 random images serve as tuning set and similarly 30 (distinct) random images are used for validation (the remaining images are not used to train the node).

6.3 Test dataset (BRATS 2016)

The pipeline described above is fully automated. To our knowledge, the BRATS 2015 training dataset pre-processing includes rigid registration (as well as resampling to a common image geometry), bias field correction and skull stripping [9]. The BRATS 2016 test dataset contains a number of unprocessed or partially pre-processed images (cf. Fig. 4). To cope with that, the pipeline was modified to include rigid registration and resampling, bias field correction [16] and skull stripping as part of a semi-automatic pre-processing step.

7 Results

The proposed approach is implemented in a .NET based DF framework which we call BONSAI. All experiments were performed on a 3.6GHz Intel Xeon processor system with 16GB RAM running Microsoft Windows 10. Training on the BRATS 2015 dataset takes around 6 to 7 hours (including "testing" on the whole dataset). Testing takes about

¹ As an illustration on WT layers. The 4 (WT) clusters are obtained from (the single DF of) the first WT layer. The second and third WT layers each consist of 4 distinct (50-tree) DFs, each of which is trained on cluster-specific data. At test time, voxels \mathbf{x} pass through the first WT layer and are assigned a cluster $k_{\mathbf{x}} \in \{1 \cdot \cdot \cdot 4\}$. Then for the second and third layers, they are sent through the DF specific to the $k_{\mathbf{x}}$ -th cluster. The same process is followed for TC and ET layers.

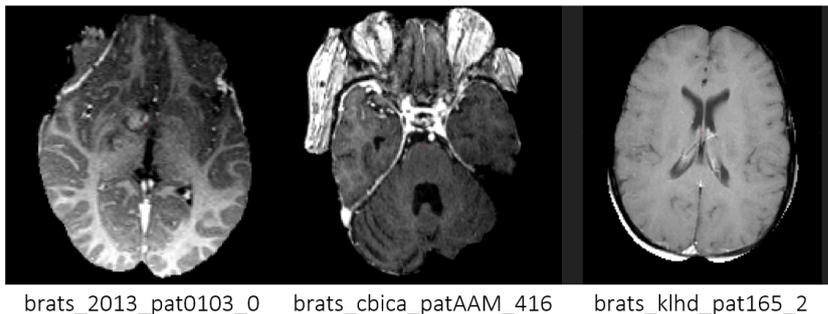


Fig. 4. BRATS 2016 test data: example of variability not seen in the training set. (Left) Bias field, (Middle) Partial skull stripping, (Right) Rigid misalignment, different geometry.

20 seconds per image. The model is trained as described previously on 70 images from the BRATS 2015 training set (about 1/4th). Table 1 reports accuracy over the remaining of the BRATS 2015 dataset. These numbers are comparable to those reported in the

Dice	WT	TC	ET
Avg \pm Std.	84% \pm 15%	72% \pm 26%	71% \pm 29%
Median (IQR)	90% (13%)	82% (28%)	82% (30%)

Table 1. Hold-out accuracy on the BRATS 2015 dataset for the model entered in the challenge *i.e.*, statistics over the remaining of the BRATS 2015 dataset after training on 70 images.

recent state-of-the-art *e.g.* [8]. It is informative to compare the predicted accuracy (on validation data) against the performance on the BRATS 2016 test set. Rough estimates of the median DICE scores are as follows (exact numbers are not available): about 80% for WT, 60% for TC and over 70% for ET. Proper delineation of the tumour core seems to remain the most challenging task, perhaps due to the high variability of appearance across images and glioma types.

8 Conclusion

We described a principled method to train DFs using hold-out estimates of the predictive error, lifting the accuracy and generalization of individual nodes and of the DF altogether. We find that shallow lifted trees formed of a few dozens or hundreds of nodes favorably compare to conventional deep trees formed of millions of nodes. This is of practical interest: it makes training, tuning and experimenting with randomized DFs much more straightforward. We exploit this benefit to experiment within the framework of auto-context forests, on challenging multi-class and multi-organ medical image segmentation tasks. We report our experience using this approach in the MICCAI 2016 BRATS challenge, where it belongs to the top performers.

Acknowledgment. The authors would like to thank the Microsoft–Inria Joint Centre for partially funding this work.

References

1. Amit, Y., Geman, D.: Shape quantization and recognition with randomized trees. *Neural Computation* 9(7), 1545–1588 (1997) [3](#)
2. Archambeau, C., Verleysen, M.: Robust Bayesian clustering. *Neural Networks* 20(1), 129–138 (2007) [7](#)
3. Breiman, L.: Bagging predictors. *Machine Learning* 24(2), 123–140 (1996) [3](#)
4. Cordier, N., Delingette, H., Ayache, N.: A patch-based approach for the segmentation of pathologies: application to glioma labelling. *IEEE T Med Imaging* 35(4) (2015) [8](#)
5. Criminisi, A., Robertson, D., Konukoglu, E., Shotton, J., Pathak, S., White, S., Siddiqui, K.: Regression forests for efficient anatomy detection and localization in computed tomography scans. *Medical Image Analysis* 17(8), 1293–1303 (2013) [3](#)
6. Geremia, E., Menze, B.H., Ayache, N.: Spatially adaptive random forests. In: 2013 IEEE 10th International Symposium on Biomedical Imaging. pp. 1344–1347. IEEE (2013) [1](#)
7. Ho, T.K.: The random subspace method for constructing decision forests. *IEEE Trans Pattern Anal Mach Intell* 20(8), 832–844 (1998) [3](#)
8. Kamnitsas, K., Ledig, C., Newcombe, V.F., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B.: Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. arXiv preprint arXiv:1603.05959 (2016) [10](#)
9. Menze, B., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., et al.: The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS). *IEEE T Med Imaging* 34(10), 1993–2024 (2015) [7, 9](#)
10. Menze, B.H., Van Leemput, K., Lashkari, D., Weber, M.A., Ayache, N., Golland, P.: A generative model for brain tumor segmentation in multi-modal images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 151–159. Springer (2010) [8](#)
11. Pereira, S., Pinto, A., Alves, V., Silva, C.A.: Deep convolutional neural networks for the segmentation of gliomas in multi-sequence MRI. In: International Workshop on Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries. pp. 131–143. Springer (2015) [1](#)
12. Quinlan, J.R.: Simplifying decision trees. *Int J Man Mach Stud* 27(3), 221–234 (1987) [5](#)
13. Shotton, J., Johnson, M., Cipolla, R.: Semantic texton forests for image categorization and segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2008. pp. 1–8. IEEE (2008) [1](#)
14. Tu, Z.: Auto-context and its application to high-level vision tasks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2008. pp. 1–8. IEEE (2008) [1](#)
15. Tu, Z., Bai, X.: Auto-context and its application to high-level vision tasks and 3D brain image segmentation. *IEEE Trans Pattern Anal Mach Intell* 32(10), 1744–1757 (2010) [1](#)
16. Tustison, N., Gee, J.: N4ITK: Nicks N3 ITK implementation for MRI bias field correction. *Insight Journal* (2009) [9](#)
17. Tustison, N.J., Shrinidhi, K., Wintermark, M., Durst, C.R., Kandel, B.M., Gee, J.C., Grossman, M.C., Avants, B.B.: Optimal symmetric multimodal templates and concatenated random forests for supervised brain tumor segmentation (simplified) with ANTsR. *Neuroinformatics* 13(2), 209–225 (2015) [1](#)

18. Tustison, N., Wintermark, M., Durst, C., Avants, B.: Ants andarboles. *Multimodal Brain Tumor Segmentation* p. 47 (2013) [1](#)
19. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*. vol. 1, pp. I–511. IEEE (2001) [3](#)
20. Zhang, Y., Brady, M., Smith, S.: Segmentation of brain MR images through a hidden Markov Random Field model and the Expectation-Maximization algorithm. *IEEE T Med Imaging* 20(1), 45–57 (2001) [8](#)
21. Zikic, D., Glocker, B., Konukoglu, E., Criminisi, A., Demiralp, C., Shotton, J., Thomas, O., Das, T., Jena, R., Price, S.: Decision forests for tissue-specific segmentation of high-grade gliomas in multi-channel MR. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 369–376. Springer (2012) [1](#)

A BRATS 2015 dataset: training IDs

For completeness, the identifiers of images from the BRATS 2015 dataset that were used for training (section [6.1](#)) are listed below.

2013_pat0001_1, 2013_pat0002_1, 2013_pat0003_1, 2013_pat0004_1, 2013_pat0005_1,
 2013_pat0006_1, 2013_pat0007_1, 2013_pat0008_1, 2013_pat0009_1, 2013_pat0010_1,
 2013_pat0011_1, 2013_pat0012_1, 2013_pat0013_1, 2013_pat0014_1, 2013_pat0015_1,
 2013_pat0022_1, 2013_pat0024_1, 2013_pat0025_1, 2013_pat0026_1, 2013_pat0027_1,
 tcia_pat105_0001, tcia_pat117_0001, tcia_pat124_0003, tcia_pat133_0001, tcia_pat149_0001,
 tcia_pat153_0181, tcia_pat165_0001, tcia_pat170_0002, tcia_pat260_0129, tcia_pat260_0244,
 tcia_pat260_0317, tcia_pat265_0001, tcia_pat290_0580, tcia_pat296_0001, tcia_pat300_0001,
 tcia_pat314_0001, tcia_pat319_0001, tcia_pat370_0001, tcia_pat372_0001, tcia_pat375_0001,
 tcia_pat377_0001, tcia_pat396_0139, tcia_pat396_0176, tcia_pat401_0001, tcia_pat430_0001,
 tcia_pat491_0001, 2013_pat0001_1, 2013_pat0004_1, 2013_pat0006_1, 2013_pat0008_1,
 2013_pat0011_1, 2013_pat0012_1, 2013_pat0013_1, 2013_pat0014_1, 2013_pat0015_1,
 tcia_pat101_0001, tcia_pat109_0001, tcia_pat141_0001, tcia_pat241_0001, tcia_pat249_0001,
 tcia_pat298_0001, tcia_pat307_0001, tcia_pat325_0001, tcia_pat346_0001, tcia_pat354_0001,
 tcia_pat393_0001, tcia_pat402_0001, tcia_pat408_0001, tcia_pat413_0001, tcia_pat442_0001,
 tcia_pat449_0001,