# Atlas Forests
# Multi-Atlas Label Propagation with Atlas Encoding by Randomized Forests

Darko Zikic, Ben Glocker, and Antonio Criminisi

Microsoft Research Cambridge

**Abstract** We describe our submission to the MICCAI 2013 SATA Challenge. The method is based on multi-atlas based label propagation, its major characteristic being that is uses the concept of an *atlas forest* to represent an atlas. This results in an efficient scheme, which requires only a single registration to label a target. Fusion of the probabilistic label proposals from each atlas is done by averaging across atlases. Results are submitted for the unregistered Diencephalon data set.

## 1 Method Overview

This submission is based on the concept of *atlas forests*, which is presented in detail in [1]. Therefore, we describe the approach only at a high-level, and list specific modifications of [1]. The submitted approach follows the standard multi-atlas label propagation model. For a given target image, a label proposal is generated by each atlas, and the proposals are then fused into a final labeling. The special characteristic of our approach is that each atlas is represented by a randomized classification forest, which is trained only on this atlas. We call this an atlas forest (AF). The approach is designed with the goal of efficiency. To label a target, only a single registration is needed, which is in contrast to most existing approaches, which require the registration of all atlases to the target. This registration aligns a probabilistic atlas to the target, and the aligned label priors are then used to augment the original input for the processing by the AFs. The actual evaluation of the atlas forests is also highly efficient. At test time, each atlas forest produces a probabilistic label estimate, which are then fused by averaging. The approach is summarized in Fig. 1.

Compared to [1], we modified the system for the challenge submission in two points, due to the comparably small number of labels. This allowed us to use context-sensitive features for label priors, now effectively treating prior and intensity channels in the same way ([1] used only local features were used on priors). Also, we did not use hierarchical priors for the challenge submission.

Finally, we tested two new variations. The first artificially augments the training set by left/right *mirroring* of each atlas. In the challenge system, submissions marked *Attempt No. 1/2* are based on the original/augmented training set.

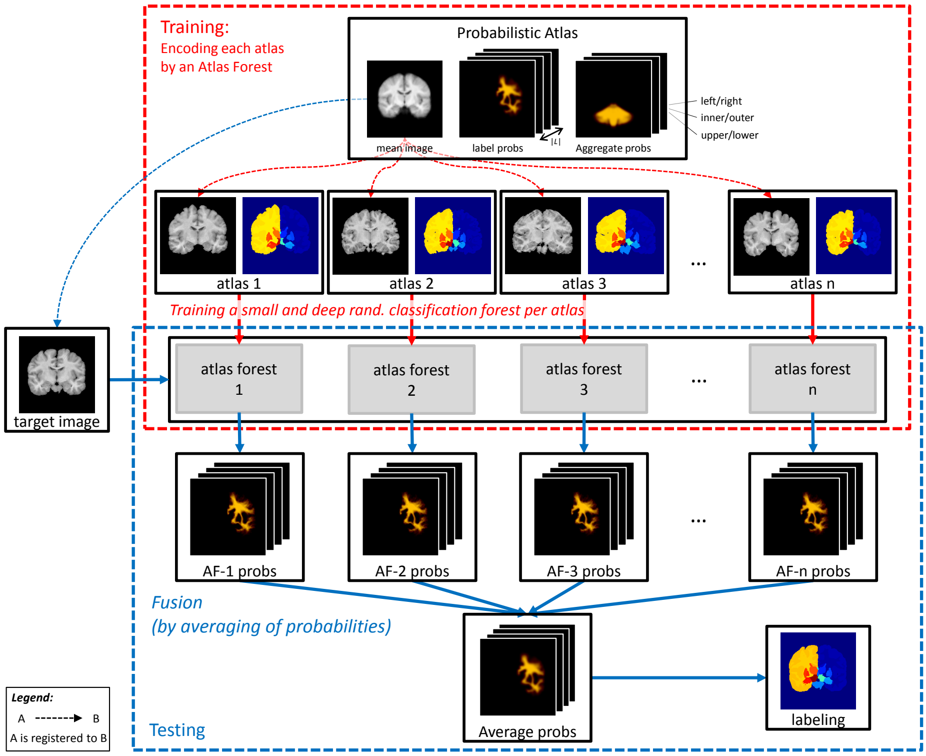


Figure 1: Overview of the atlas forest framework.

The second variation (*Stage 2*) consists of running the standard AF framework, and then using its probabilistic output to augment the target image (instead of the aligned probabilistic atlas priors), before running AFs on this new input in a 2nd stage. The motivation is to remove dependency on the registration scheme. The 2nd stage AF is specifically trained on the new version of the augmented input data. Performing this training properly (i.e. excluding the target in training of classifier in stage 1 in the cross-validation setting) is much easier to achieve with the AF framework than with standard forest schemes.

## 2 Experiments

We submitted our results for the Diencephalon data set. Since we do not require registrations of all atlases to the target, we used only the unregistered versions of the data set. Here, we describe the settings of the system and training, and summarize the challenge results in Tab. 2 – for details please see the challenge web page `http://masi.vuse.vanderbilt.edu/submission/leaderboard.html`. Additionally to the AF variants, we report baseline results for a subcomponent, which transfers labels to the target directly from the aligned probabilistic atlas.

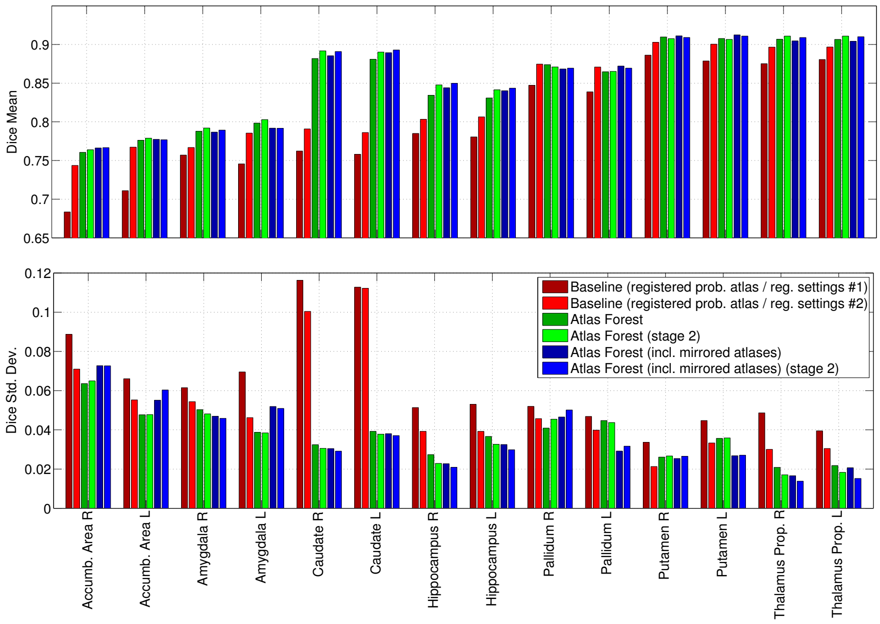


Figure 2: Results of leave-one cross-validation on training data set.

*Preprocessing* The input images are cropped, such that the brain is centered. We perform inhomogeneity correction by [2], and histogram equalization with the first image in the training data set as reference. No skull stripping is performed.

*Generation of the Probabilistic Atlas and Registration Settings* The probabilistic atlas is constructed by iterative non-linear registration of all training images to their mean, with three iterations. This results in an average intensity image, and a set of label priors, which are additionally smoothed by a Gaussian with $\sigma = 2$mm. We use affine, followed by deformable FFD-based registration [3], with cross-correlation as data term, strong regularization, and control point spacing of 30mm. An image pyramid with 4 levels is employed. We tested two different registration settings, the first (used in [1]) operates only on the two coarsest pyramid levels, while the second uses the three coarsest 3 levels. The second setting significantly improves the results (cf. Fig. 2, Tab. 1) and was used for all reported experiments. The same registration is used for generation of the probabilistic atlas, and its alignment to the targets.

*Forest Settings* Method setup was done by leave-one-out cross-validation on the training data, while monitoring the Dice score. The final results of these experiments are summarized in Fig. 2 and Tab. 1. We use 5 trees per atlas forest, with maximal depth of 36, and the restriction that each leaf contains at least 8 samples. Each node uses a randomly chosen batch of 2000 random features, with 10 different batches available per tree. The same settings are used for all reported results.

| Method | Dice mean | Dice std. | Dice median |
|---|---|---|---|
| Baseline (reg. settings #1) | 0.7992 | 0.0622 | 0.8149 |
| Baseline (reg. settings #2) | 0.8280 | 0.0506 | 0.8407 |
| Atlas Forest | 0.8514 | 0.0370 | 0.8556 |
| Atlas Forest - S2 | 0.8557 | 0.0359 | 0.8602 |
| Atlas Forest (mirrored) | 0.8538 | 0.0363 | 0.8588 |
| Atlas Forest - S2 (mirrored) | 0.8556 | 0.0359 | 0.8610 |

Table 1: Summary of leave-one-out cross-validation on training data.

| | Dice | | | Hausdorff | | |
|---|---|---|---|---|---|---|
| Method | mean | std. | median | mean | std | median |
| Baseline (reg. settings #1) | 0.7980 | 0.0493 | 0.8105 | 4.1327 | 0.7732 | 3.8543 |
| Atlas Forest | 0.8208 | 0.0544 | 0.8452 | 3.9868 | 0.9303 | 3.7594 |
| Atlas Forest - S2 | 0.8237 | 0.0566 | 0.8462 | 3.8292 | 0.9351 | 3.6449 |
| Atlas Forest (mirrored) | 0.8248 | 0.0470 | 0.8422 | 3.7411 | 0.7833 | 3.5231 |
| Atlas Forest - S2 (mirrored) | 0.8282 | 0.0495 | 0.8484 | 3.8659 | 0.8387 | 3.6737 |

Table 2: Summary of results on testing data from the challenge system.

*Runtime* The training of one tree takes 15-20 minutes on a standard PC. After pre-processing the target image, and aligning the probabilistic atlas to it, the label propagation takes less than 2 minutes. For the alignment of the probabilistic atlas to the target image, the registration with the used settings (#2) takes less than 2 minutes per image, and the warping of the label maps less than 1 minute.

## 3 Concluding Remarks

On the challenge task, we do not quite achieve the accuracy of the best submissions, with a mean Dice difference of up to 0.04. Interestingly, the results are en par with *Baseline Majority Vote* on the registered data, which can be seen as a corresponding approach using traditional atlas representation. Given its simplicity, we consider the submission to show the potential of atlas forests as a basis for efficient and accurate labeling schemes. Further improvement might be achieved through atlas selection, and use of advanced fusion and classification.

## References

1. Zikic, D., Glocker, B., Criminisi, A.: Atlas encoding by randomized forests for efficient label propagation. In: MICCAI. (2013)
2. Tustison, N., Gee, J.: N4ITK: Nick's N3 ITK implementation for MRI bias field correction. The Insight Journal (2010)
3. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through MRFs and efficient linear programming. MedIA (2008)