

# Iterative Next Boundary Detection for Instance Segmentation of Tree Rings in Microscopy Images of Shrub Cross Sections

Alexander Gillert<sup>1</sup> Giulia Resente<sup>2</sup> Alba Anadon-Rosell<sup>3</sup>  
 Martin Wilmking<sup>2</sup> Uwe Freiherr von Lukas<sup>1,4</sup>

<sup>1</sup>Fraunhofer Institute for Computer Graphics Research (IGD), Rostock

<sup>2</sup>Institute of Botany and Landscape Ecology, Ernst Moritz Arndt University, Greifswald

<sup>3</sup>Centre for Research on Ecology and Forestry Applications (CREAF), Barcelona

<sup>4</sup>Institute for Visual & Analytic Computing, University of Rostock

{alexander.gillert, uwe.freiherr.von.lukas}@igd-r.fraunhofer.de

## Abstract

We address the problem of detecting tree rings in microscopy images of shrub cross sections. This can be regarded as a special case of the instance segmentation task with several unique challenges such as the concentric circular ring shape of the objects and high precision requirements that result in inadequate performance of existing methods.

We propose a new iterative method which we term *Iterative Next Boundary Detection (INBD)*. It intuitively models the natural growth direction, starting from the center of the shrub cross section and detecting the next ring boundary in each iteration step. In our experiments, INBD shows superior performance to generic instance segmentation methods and is the only one with a built-in notion of chronological order.

Our dataset and source code are available at <http://github.com/alexander-g/INBD>.

## 1. Introduction

Dendrochronology is the science that provides methodologies to date tree rings [4], i.e. measuring and assigning calendar years to the growth rings present in a wood stem. By analyzing anatomical properties like ring widths or the cell sizes within the rings, dendrochronology can be applied to dating archaeological manufactures, tracking timber sources or reconstructing past climate conditions [11].

For climate reconstruction in the Arctic, shrubs constitute the most important source of dendrochronological information, since they are the only woody plants able to thrive there [23]. As temperature is a limiting factor for shrub growth in the Arctic, it shows a strong relationship




Figure 1. Example microscopy images (left) of shrub cross sections from our new dataset and the outputs (right) of our proposed method INBD for instance segmentation of tree rings

with climate, making these plants a reliable proxy to reconstruct past climate events [24]. Dendrochronological analyses on shrubs are usually performed on thin cross sections of branches or roots and observed under the microscope with a magnification that allows ring identification at a cellular level. As of now, ecological studies are limited in size by the amount of manual analysis work due to the lack of automatic tree ring detection methods.

With this paper we want to introduce this problem to the

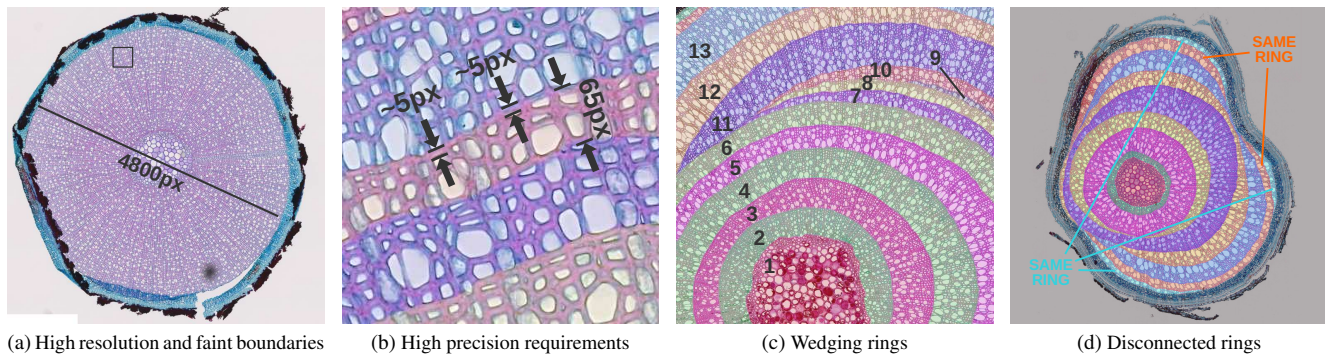


Figure 2. Some of the challenges encountered in this task:

- (a) Boundaries inbetween tree rings are often hard to recognize. For example, this cross section contains 14 rings.
- (b) Crop of the previous image (indicated by the square) with overlaid annotation. A tree ring is only 65 pixels wide or ca 1.4% of the full cross section diameter. The cell wall that divides late summer cells and the next year’s early summer cells is only 5 pixels wide or 0.1%.
- (c) Wedging rings can complicate finding the chronologically correct next year ring.
- (d) Rings can grow in multiple disconnected parts from different sides.

computer vision community and enhance the capabilities for ecological sciences. We release a new dataset containing high resolution microscopy images of shrub cross sections and propose a specialized method for growth ring identification. Example images from our dataset and corresponding outputs of our method are shown in Figure 1. From a computer vision point of view, this can be regarded as a special case of the instance segmentation task, however it differs from previous generic datasets in several ways which makes existing methods underperform.

Figure 2 illustrates these differences. For one, the concentric ring shape of the instances can pose a significant obstacle, particularly for top-down methods because the objects have almost identical bounding boxes. This gets complicated by the fact that year rings can also form incomplete circles (wedging rings) and grow from only one side, or even in multiple disconnected parts from different sides (2d). Depending on the species, plant part and climatic conditions the amount of wedging rings can range from zero to being the majority. Assigning the correct order to wedging rings can be an issue where rings of more than 2 years touch each other (2c). Bottom-up methods on the other hand struggle with faint ring boundaries (2a) as the presence of the boundary pattern is not always constant throughout the whole stem circumference. They are prone to merging rings where no boundary can be detected or splitting them where the ring width is narrow. Next, the images are acquired at a high resolution (2a) to capture cellular information, yet a high degree of precision is required for the downstream task of assigning individual cells to the correct year. The thickness of a cell wall that is dividing the cells from one ring to another can be as low as 0.01% of the whole object (2b). Finally, as the preparation of samples and annotation of the

images is very costly, training has to be performed in a low data regime.

We argue that a specialized approach can help to overcome those challenges and propose a new iterative method which we term Iterative Next Boundary Detection (INBD). In the first step, it performs semantic segmentation to detect basic features such as the background, center and the ring boundary pixels. From this starting point, it iteratively detects the next year ring’s boundaries, following the natural growth of the plant. This process is augmented with a recurrent wedging ring detection module to counteract issues with incomplete rings. We compare our method with both top-down and bottom-up generic instance segmentation in our experiments in which it shows better results. Moreover, it is the first method that automatically assigns a chronological order to the detected objects.

The contributions of this paper can summarized as follows:

- Publication of a new challenging dataset for a special case of instance segmentation.
- Development of the specialized method INBD for tree ring instance segmentation.
- Evaluation of previous generic instance segmentation methods and comparison with INBD

## 2. Related Work

Instance segmentation is a widely studied problem in computer vision, commonly benchmarked on a variety of standard generic datasets such as COCO [12] which contains photographs of everyday objects or the the more specialized CREMI 2016 [6] challenge for cell segmentation in

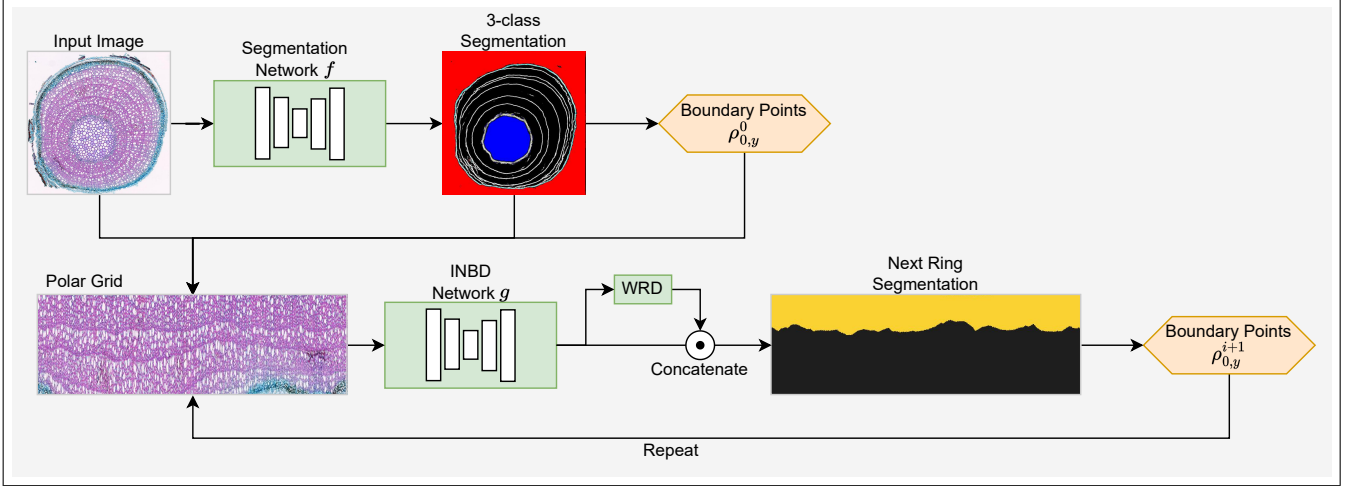


Figure 3. Overview of the INBD pipeline. An input image is first passed through a generic semantic segmentation network that detects 3 classes: background, ring boundaries and the center ring. A polar grid is sampled starting from the the detected center ring and passed to the main INBD network that detects the next ring. This process is repeated until the background is encountered.

electron microscopy images. No publicly available dataset is known to us that contains concentric ring shaped and ordered objects.

Methods can be categorized into *top-down* and *bottom-up* procedures. Mask-RCNN [8] is the most widely used architecture and belongs to the top-down group. It relies on an object detector to first detect bounding boxes of objects which are then segmented. This fails on overlapping or as in our case concentric objects due to non-maximum suppression. Moreover, it can only generate low resolution masks. Contour methods such as Deep Snake [17] or DANCE [13] can generate masks with higher precision but still require an upstream object detector.

Bottom-up methods for instance segmentation methods work by first computing object boundaries or affinities and then clustering the resulting superpixel graph into whole objects via the multicut objective. Finding the optimal solution for this is known to be NP-hard [10], therefore several approximate solvers such as GASP [2] have been developed. These methods perform significantly better on our dataset but still show deficits in cases where object boundaries are hard to recognize and they cannot handle disconnected rings (such as in Fig. 2d).

None of the above methods has a built-in notion of sequence order of the detected objects that would be needed to assign a tree ring to a year.

Application of deep learning methods to ecological purposes is nowadays an established procedure [3] due to the complexity related to ecological investigations and the use of increasingly larger datasets. Specifically for quantitative wood anatomy (QWA), deep learning research has so far focused mostly on detection and measurement of cells such as in [7, 19]. Tree ring detection was subject in [5, 18], however

only on scans or photographs of mature wood core samples rather than full cross sections as in our case.

ROXAS [21, 22] is the most commonly used analysis tool in QWA, however it is based on traditional image processing methods and not on deep learning which makes it sensitive to sample processing and image quality. It also contains tree ring detection functionality which works by line-following early summer cells but requires domain knowledge for manual tuning of many species-specific parameters like cell shape and size.


### 3. Method

On a high level, INBD simply modifies and extends the various contour based methods like Deep Snake [17] or PolarMask [26] with an iterative inference procedure. In reality, this requires several important design choices to make this perform well. The influence of the individual design choices is analyzed in an ablation study in subsection 5.2. An architectural overview of the INBD pipeline can be found in Figure 3.

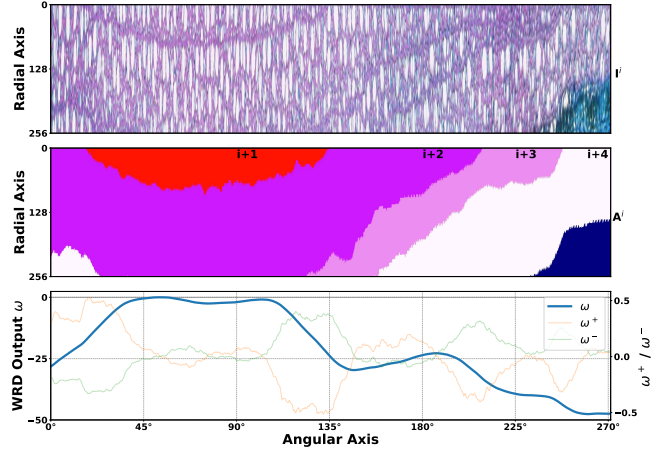
#### 3.1. Network Architecture

The INBD pipeline is composed of two neural networks. The first network is a simple semantic segmentation network that is trained to detect three classes: background, ring boundaries and the center ring (or pith). We denote this network and its output with  $f(I) = (y^{bg}, y^{bd}, y^{ct})$  when applied on image  $I$ . We select an architecture based on U-Net [20] with a pretrained backbone. The three classes are trained with a combination of cross-entropy loss and the Dice loss [16]:

$$L_f = \lambda_1 L_{CE}^{background} + \lambda_2 L_{Dice}^{boundaries} + \lambda_3 L_{CE}^{center} \quad (1)$$



(a) Polar Grid Sampling



(b) Wedging Ring Detection

Figure 4. Visual explanation of the concepts. (a) We sample on a polar grid starting from the previous detected ring boundary. The number of points is reduced for better visualization. (b) shows the resulting input to the network  $g$  (top), the corresponding annotation (center) and the output signal  $\omega$  of the wedging ring detection module (bottom).  $\omega$  accumulates along the angular axis, rising on start and falling on end of wedging rings.

with  $\lambda = (0.01, 1.0, 0.1)$  balancing coefficients to account for class imbalances. Due to the large size of the images and for a larger field of view,  $f$  operates on  $\times 0.25$  of the original resolution.

Our main network, which we denote with  $g$ , is another 2D convolutional segmentation network that classifies each pixel as belonging to the next ring or not. By choosing a 2D network as opposed to a 1D one, as in many contour methods such as Deep Snake [17], we can leverage transfer learning since we are working in a low data regime and in addition to that we can reject and interpolate ambiguous predictions (see below, Eq. 4). This second network has mostly the same architecture as the first one, except that we replace the normal 2D convolutions with circular convolutions to wrap around the full circle, as also used in Deep Snake [17]. The circularity is only applied to the angular axis (see below).

### 3.2. Polar Grid

The network  $g$  operates on “unrolled” rings  $I^i \in \mathbb{R}^{[C \times N \times M]}$  sampled on a polar grid  $P^i \in \mathbb{R}^{[N \times M]}$  with  $N = 256$  a fixed resolution in the radial dimension and  $M$  an adaptive resolution in the angular dimension. The polar grid origin is computed from the center of mass of the center ring  $y^{ct}$  as detected by  $f$ . Polar coordinates impose a prior, ensuring a coherent (quasi-)convex shape, contrary to Cartesian coordinates.

We express the sampling points for ring  $i$  as polar coordinates  $(\rho_{xy}^i, \varphi_{xy}^i)$ , with  $x \in [0, N]$ ,  $y \in [0, M]$  indices within the grid. The boundary point radii for the second ring  $\rho_{0,y}^1$  are inferred directly from the detected center ring  $y^{ct}$ .

Estimating the extent of the grid in the radial dimension (i.e.  $\rho_{N,y}$ ) is crucial: if too short, the next ring will be cut off, if too long, the next ring might get skipped and not detected at all. For this, we compute the distances to the closest positive value in  $y^{bd}$  for each angle  $\varphi$  and set the extent to  $1.5 \times 95\%$ -th percentile of these distances, to make sure that most points are included but also to filter outliers. This was empirically verified to cover all rings in our dataset. The remaining radial values  $\rho$  are then uniformly distributed along this range:  $\rho_{x,y} = \frac{1}{N}(\rho_{N,y} - \rho_{0,y})x + \rho_{0,y}$

The angular resolution  $M$  is computed so that the angles  $\varphi$  have an approximately uniform euclidean distance to each other across rings: since the outer rings have a larger circumference than the inner ones they should be sampled at a higher angular resolution  $M$ . The value  $M^i$  for ring  $i$  is computed from the previous ring’s average radii:

$$M^i = \alpha \frac{1}{M^{i-1}} \sum_y^{M^{i-1}} \rho_{0,y}^{i-1} \quad (2)$$

with  $\alpha$  a hyperparameter that controls the general density of  $M^i$  which we set to  $2\pi$  where not otherwise mentioned. The angles  $\varphi$  are spaced uniformly along the full circle:  $\varphi_{x,y} = \frac{1}{M} 2\pi y$

The channel dimension  $C = 7$  is composed of the RGB channels of the input image, the detected “background” and the “boundaries” outputs  $y^{bg}$  and  $y^{bd}$  from  $f$ , normalized radii  $\rho$  and the output of the wedging ring detection module (see below) concatenated together.

The main loss  $L_{CE}^{cls}$  for network  $g$  is the standard cross-entropy loss to classify each pixel in the polar grid as belonging to the next ring or not, according to the annotation  $A^i$ , sampled on the same polar grid.

### 3.3. Inference

To perform inference of the next ring's boundary points  $\rho_{0,y}^{i+1}$ , we select last positive point in the output  $g(I^i)$  column-wise, where it is unambiguous:

$$\tilde{X}_y = \{x, \text{ where } g(I^i)_{x,y} = 1\} \quad (3)$$

$$\rho_{0,y}^{i+1} = \begin{cases} \rho_{\max \tilde{X}_y, y}^i, & \text{if } \max \tilde{X}_y = \min \overline{\tilde{X}_y} - 1 \\ \text{undefined}, & \text{otherwise} \end{cases} \quad (4)$$

Ambiguous values linearly interpolated. Importantly, the interpolation should be performed on polar coordinates and not on Cartesian ones and wrap around the circle. This detection process is repeated iteratively with the new predicted ring boundary points  $i + 1$  as the starting point to detect the ring  $i + 2$  until the background  $y^{bg}$  that was detected by the segmentation network  $f$  is reached.

### 3.4. Wedging Ring Detection

The method as described so far is able to detect full tree rings sufficiently well but struggles with wedging rings. More specifically, it is prone to skipping a ring boundary in locations where the wedging ring is far away and outside the field of view, e.g. as in Figure 2c when trying to detect the next boundary after ring 6. To counteract this issue we insert a wedging ring detection (WRD) module before the final classification layer.

This module consists of 3 additional convolutional layers with two output channels. The two channels are averaged along the radial axis into 1-dimensional signals  $\omega^+$  and  $\omega^- \in \mathbb{R}^M$  and combined via a recurrent mechanism:

$$\omega'_0 = \beta \quad (5)$$

$$\omega'_\varphi = \sigma(\omega_{\varphi-1}^+) - \sigma(\omega_{\varphi-1}^-) \quad (6)$$

$$\omega_\varphi = \omega'_\varphi - \max \omega' \quad (7)$$

where  $\sigma$  is the sigmoid function and  $\beta$  is a starting point constant. Intuitively,  $\omega^+$  is responsible for detecting the start of a wedging ring and increases the output signal  $\omega$ , whereas  $\omega^-$  detects the end and decreases it.  $\omega$  is then forwarded to the final classification layer by concatenating it to the features along the channel dimension.

During inference, the choice of  $\beta$  does not matter because of the normalization by subtracting the maximum (Eq. 7). This ensures a standardized representation of the signal to the following downstream classification layer, irrespective of the starting point  $\beta$ . High values close to zero indicate valid locations (next ring or  $i + 1$ ), whereas low values are invalid locations (next but one or  $i + 2$ ). This functionality is illustrated graphically in Figure 4b.

Although in theory the network could derive useful information from this module by itself, we have found that in practice it is highly beneficial to add an explicit training

signal. Again, we use the cross entropy loss, but modified for the single dimension and applied on the unnormalized signal  $\omega'$ :

$$L^{wrd} = A_\varphi^{wrd} \log \sigma(\omega'_\varphi) + (1 - A_\varphi^{wrd}) \log 1 - \sigma(\omega'_\varphi) \quad (8)$$

$$A_\varphi^{wrd} = \begin{cases} 1, & \text{where } A_{0,\varphi}^i = i + 1 \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

During training  $\beta$  is set so that  $\sigma(\omega') = 0$  if the ground truth at angle  $\varphi = 0$  is low, or so that  $\sigma(\omega') = 1$  if it is high, to avoid incorrect training signals (we choose  $\beta = \pm 15$ ).

An example where this module helps to catch an error is shown in Figure 5.




Figure 5. INBD can be prone to skipping boundaries. In this example, the wedging ring detection module helps to catch mistakes like this. (Left: without WRD, right: with WRD)

The final loss for network  $g$  is defined as  $L_g = L_{CE}^{cls} + \lambda_4 L^{wrd}$  with  $\lambda_4 = 0.01$ . We have found higher values to have a negative impact on the main classification loss.

### 3.5. Training Procedure

Since INBD is an iterative procedure, errors caused by an earlier ring get easily propagated onto the later rings. It can however also recover from previous mistakes if trained with an *iterative training* procedure: rather than using only the (near-perfect) boundary points from the annotation, the training loop should incorporate previous (possibly faulty) predictions as the starting point for polar grids. Listing 1 shows the high-level pseudo code for one training epoch.

Listing 1. Pseudo-code for one training epoch

```

for (image  $I$ , annotation  $A$ , ring  $i$ ) in dataset:
     $L_g = 0$ 
     $\rho_{0,y}^i = \text{boundary\_from\_annotation}(A, i)$ 
    loop  $i = i..i+n$ :
         $\hat{\rho}_{0,y}^i = \text{augment}(\rho_{0,y}^i)$ 
         $I^i = \text{sample\_polar\_grid}(I, (\hat{\rho}^i, \varphi^i))$ 
         $y^i = g(I^i)$ 
         $L_g += \text{compute\_loss}(y^i, A^i)$ 
         $\rho_{0,y}^{i+1} = \text{compute\_boundary}(y^i)$ 
    backpropagate(  $L_g / n$  )

```

Subset	Species	Training Images	Test Images	Number of rings	Average diameter	Example Images
DO	Dryas octopetala	22	42	544	3700px	Figs. 1 (bottom), 7 (bottom)
EH	Empetrum hermaphroditum	24	58	949	3260px	Figs. 1 (top), 2c, 2d
VM	Vaccinium myrtillus	22	45	494	3979px	Figs. 2a, 2b, 7 (top)

Table 1. Overview of our dataset

Where not otherwise mentioned we use  $n = 3$  iterations per epoch.

Besides the conventional data augmentations such as the pixel-wise color jitter operations we employ additional augmentations specific to polar grids such as varying the boundary points:

$$\hat{\rho}_{0,y} = \rho_{0,y} + \cos(\varphi_{0,y} + X_0)\gamma_0 + X_1\gamma_1 \quad (10)$$

with  $X \sim \mathcal{U}(-1, 1)$  random variables and  $\gamma$  hyperparameters.

Both networks are trained separately with the AdamW [15] optimizer for 100 epochs,  $1e-3$  base learning rate and cosine annealing [14] learning rate schedule.

## 4. Experimental Setup

### 4.1. Dataset

Our dataset consists of overall 213 high-resolution images. It is split into 3 subsets according to the plant species. An overview is provided in Table 1. The amount of images is rather low due to the high cost of sample preparation as well as annotation: a single image containing a large amount of rings can take up to 6 hours to annotate by hand. The dataset and annotations are publicly available at <http://github.com/alexander-g/INBD>.

The shrub samples were collected at subalpine, alpine and subarctic sites across the Pyrenees, Southern Norway and Northern Sweden. Aboveground shoots (ramets) were clipped at the stem base, above the soil surface. In the lab, the samples were cut into 15-20  $\mu m$  cross-sections with a rotary microtome, stained with a mixture of 1:1 safranin and astrablue, rinsed with ethanol solutions, embedded in Euparal, dried and finally scanned in a slide scanner to obtain high resolution images.

### 4.2. Compared Methods

As there are no specialized methods for tree ring detection in shrub cross sections yet, we compare our method with generic instance segmentation methods. From the top-down category we compare with Mask-R-CNN [8] and Deep Snake [17]. Mask-R-CNN is trained in two modes: in the *hollow* (h) mode, objects are defined as single calendar years and are donut-shaped (with a hole), whereas in the

*filled* (f) mode, objects consist of multiple years (and have no holes). We use the implementation from the torchvision (v0.11) framework. The non-maximum suppression is increased to 0.7 to reduce the filtering of overlapping detections and the images are downscaled to accommodate for GPU memory limits. For Deep Snake only the filled mode is used because it cannot model hollow objects.

In the bottom-up group we select Multicut [10] and GASP [2] for comparison. We use the implementation from the PlantSeg [25] source code which was developed in part by the original GASP algorithm authors. For a fair comparison, the detected boundaries from the same segmentation network  $f$  as for INBD are used. We have found bottom-up methods require species-specific tuning of hyperparameters. We have tested several combinations and report only the best ones here. More information can be found in the supplement.

### 4.3. Metrics

Our main evaluation metric is the mean Average Recall (mAR) as defined in [9] averaged at IoU=.50:.05:.95 intervals. We do not use the mean Average Precision (mAP) that is often used in generic instance segmentation literature, as we regard instance recall as more important than precision: it is easier for the end user to delete false positive objects on manual inspection than adding new ones.

We additionally report the Adapted Rand errors (ARAND) as defined in [1] because this metric is more commonly used in the bottom-up literature. It can be interpreted as the harmonic mean of the pixelwise precision and recall values.

## 5. Results

### 5.1. Method Comparison

The main results of the compared methods are presented in Table 2. For all metrics we observe consistently better performance of INBD over the compared methods.

Top-down methods show very unsatisfactory performance. The filled mode gives a small performance boost but the results are still too inaccurate to be useful, particularly missing many thin rings. Deep Snake struggles remarkably, often detecting only one or two rings at most.

Method	mAR $\uparrow$			ARAND $\downarrow$		
	DO	EH	VM	DO	EH	VM
Mask-R-CNN (h)	.106 (.008)	.144 (.003)	.185 (.008)	.644 (.007)	.694 (.002)	.532 (.004)
Mask-R-CNN (f)	.210 (.006)	.176 (.004)	.218 (.002)	.441 (.002)	.499 (.001)	.425 (.007)
Deep Snake (f)	.061 (.011)	.015 (.001)	.019 (.008)	.524 (.024)	.620 (.003)	.584 (.027)
GASP	.374 (.002)	.667 (.004)	.576 (.014)	.313 (.003)	.144 (.003)	.168 (.010)
Multicut	.387 (.008)	.688 (.005)	.596 (.006)	.301 (.001)	.132 (.004)	.154 (.005)
INBD (ours)	<b>.553</b> (.011)	<b>.738</b> (.018)	<b>.704</b> (.014)	<b>.196</b> (.009)	<b>.113</b> (.010)	<b>.112</b> (.007)

Table 2. Method comparison. Values are averaged over 3 full training runs with the standard deviation provided in parentheses. (h) refers to the hollow mode, (f) to the filled mode.  $\uparrow$  denotes higher is better,  $\downarrow$  lower is better.

We attribute this to its base detector CenterNet [27] which inherently fails with concentric objects.

The bottom-up methods can compete with INBD on EH thanks to relatively well recognizable ring boundaries in this subset. The VM and especially DO subsets on the other hand have much less pronounced and sometimes ambiguous boundaries which often cannot be detected at all. This is particularly a problem for the bottom-up methods which are then prone to incorrectly merging two rings. INBD on the other hand can interpolate ambiguous locations (Eq. 3). The results of GASP and Multicut are very similar to each other, as also noted in [25].

In general, we observe that INBD is better at detecting difficult rings. This observation is confirmed in the more fine-grained analysis in Figure 6 which shows the recall values for the individual IoU thresholds. INBD scores only slightly better on the high threshold recalls such as AR90 or AR95 which are usually the easily recognizable rings. The real benefits come from detecting harder examples.




Figure 6. Recall over IoU for the compared methods

## 5.2. Ablation Study

In Table 3 we show how the individual design choices as proposed in section 3 affect the detection performance of INBD. Two baselines of INBD without the adaptations are evaluated, one with Cartesian and another with polar coordinates. We note that those two implementations are not fully comparable, more details in the supplement.

Our experiments show that increasing the angular resolution (Eq. 2) in order to keep the Cartesian resolu-

tion roughly constant across rings yields almost a 3 mAR percentage points improvement. Interpolating ambiguous boundary points (Eq. 4) is highly important and improves the mAR by more than 6 points. Iterative training (subsection 3.5), i.e. training with previously predicted boundary points (as apposed to only using the annotation) gives an additional performance boost of more than 3 mAR. Finally, the WRD module (subsection 3.4) helps with wedging rings. As wedging rings are comparatively few in numbers, the performance gain is relatively moderate but consistent among training runs.

Configuration	mAR $\uparrow$	ARAND $\downarrow$
Cartesian coordinates baseline	.498	.237
Polar coordinates baseline	.601	.218
+ adaptive angular resolution M	.629	.190
+ ambiguous boundary interpolation	.691	.146
+ iterative training	.722	.126
+ wedging ring detection	.738	.113

Table 3. Influence of design choices on the performance. All values refer to the EH subset.

Additional evaluations on the effect of hyperparameters on the detection performance can be found in the supplementary materials.

## 5.3. Cross-species Performance

Dendro-ecological studies are rarely limited to the three plant species from our dataset, end users might want to analyze new species, for which trained models are not yet available. Therefore we test how well the compared methods generalize to unseen species. The results are presented in Table 4.

EH and VM show some level of similarity to each other and methods trained on one set can be used to a limited degree on the other one. These results might be insufficient for downstream tasks but could be used to generate new annotations for retraining, faster than creating them manually from scratch. DO on the other hand is visually dissimilar and requires networks specially trained on it.

Among the methods we observe no clear winner, though

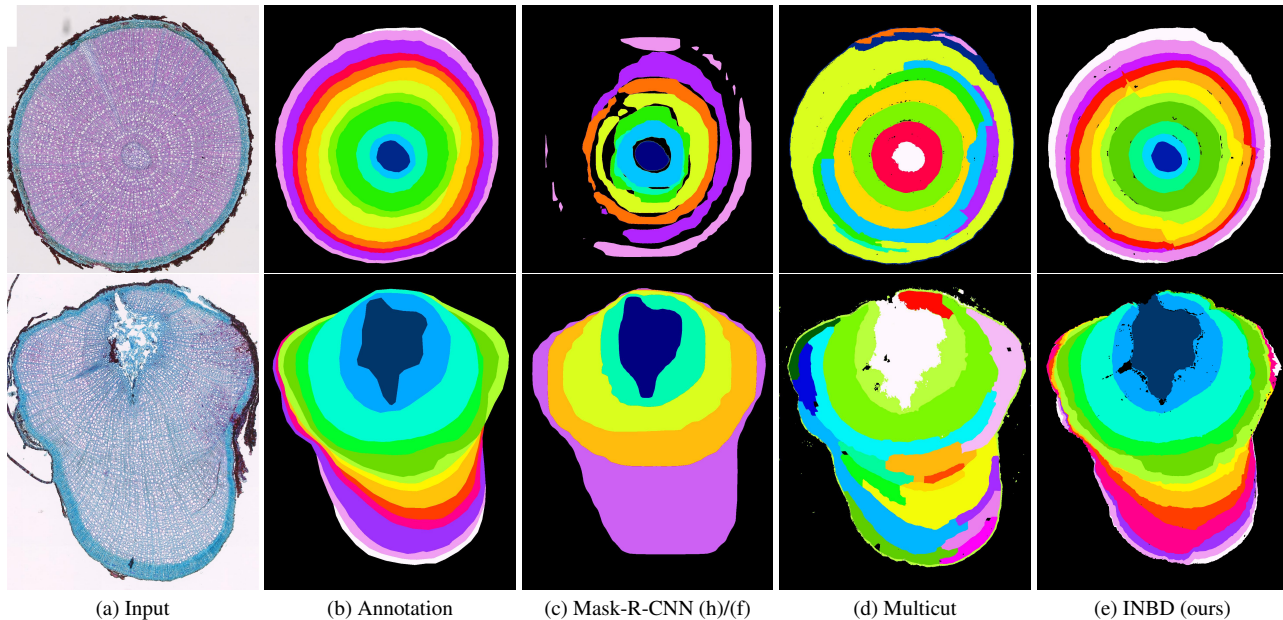


Figure 7. Qualitative comparison and examples of typical mistakes made by the compared methods

Method	Training set	Test set	mAR $\uparrow$	ARAND $\downarrow$
INBD	EH	VM	.588	.194
Multicut	EH	VM	.580	.166
INBD	VM	EH	.472	.262
Multicut	VM	EH	.393	.287
INBD	EH	DO	.106	.561
Multicut	EH	DO	.116	.552
INBD	DO	EH	.219	.435
Multicut	DO	EH	.169	.478

Table 4. Cross species ring detection performance

INBD is scoring on average slightly better. The results show that more research needs to be done into this direction.

#### 5.4. Qualitative Results

Figure 7 shows typical mistakes caused by our method as well as the compared top-down and bottom-up procedures.

INBD tends to skip boundaries and this mistake often gets propagated onto the following rings since it is an iterative procedure. However, thanks to its iterative training procedure and boundary augmentations it can still recover from this.

As expected, the detector-based Mask-R-CNN struggles with the large overlap and fails to detect many rings, and the ones that get detected are very inaccurate. Bottom-up methods such as Multicut are prone to merging rings where boundaries are difficult to recognize and to splitting them on false positive boundary detections.

More qualitative results can be found in the supplement.

## 6. Concluding Remarks

Our dataset contains only images for which annotators were confident that they are annotated correctly. In real-world ecological studies, shrub samples, especially those from harsh climatic conditions, often contain many irregularities in their anatomical structures and may be extremely difficult to fully annotate, even for experts. In addition, fully annotating images with a large number of rings is very time consuming and costly. Therefore, future work could focus on weakly supervised training from partially annotated images and on developing methods that provide a confidence estimate for each detected ring or parts of it.

Moreover, as cross section images can vary widely depending on a variety of factors such as plant species, climatic conditions or sample preparation it is not unlikely that a single method trained on a single dataset will not suffice to cover all scenarios. Further research could be performed on cross-species training for better out-of-distribution generalization.

## ACKNOWLEDGEMENTS

This work has been supported by the European Social Fund (ESF) and the Ministry of Education, Science and Culture of Mecklenburg-Vorpommern, Germany under the project "DigIT!" (ESF/14-BM-A55-0015/19).

AAR was funded by a Postdoctoral Research Fellowship from the Alexander von Humboldt Foundation (Germany) and a Juan de la Cierva-Incorporación Grant by the Government of Spain.



## References

- [1] Ignacio Arganda-Carreras, Srinivas C Turaga, Daniel R Berger, Dan Cireşan, Alessandro Giusti, Luca M Gambardella, Jürgen Schmidhuber, Dmitry Laptev, Sarvesh Dwivedi, Joachim M Buhmann, et al. Crowdsourcing the creation of image segmentation algorithms for connectomics. *Frontiers in neuroanatomy*, page 142, 2015. 6
- [2] Alberto Bailoni, Constantin Pape, Nathan Hütsch, Steffen Wolf, Thorsten Beier, Anna Kreshuk, and Fred A. Hamprecht. Gasp, a generalized framework for agglomerative clustering of signed graphs and its application to instance segmentation. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11635–11645, 2022. 3, 6
- [3] Sylvain Christin, Éric Hervet, and Nicolas Lecomte. Applications for deep learning in ecology. *Methods in Ecology and Evolution*, 10(10):1632–1644, 2019. 3
- [4] Edward R. Cook. Methods of dendrochronology. 1990. 1
- [5] Anna Fabijańska and Małgorzata Danek. Deepdendro—a tree rings detector based on a deep convolutional neural network. *Computers and electronics in agriculture*, 150:353–363, 2018. 3
- [6] Jan Funke, Stephan Saalfeld, Davi Bock, Srini Turaga, and Eric Perlman. Creml challenge. <https://cremi.org>, 2016. Accessed: 2022-10-20. 2
- [7] Ángel M. García-Pedrero, Ana I. García-Cervigón, José Miguel Olano, Miguel García-Hidalgo, Mario Lillo-Saavedra, Consuelo Gonzalo-Martín, Cristina Caetano, and Saul Calderon-Ramirez. Convolutional neural networks for segmenting xylem vessels in stained cross-sectional images. *Neural Computing and Applications*, pages 1 – 13, 2019. 3
- [8] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 3, 6
- [9] Jan Hosang, Rodrigo Benenson, Piotr Dollár, and Bernt Schiele. What makes for effective detection proposals? *IEEE transactions on pattern analysis and machine intelligence*, 38(4):814–830, 2015. 6
- [10] Jörg Hendrik Kappes, Markus Speth, Björn Andres, Gerhard Reinelt, and Christoph Schnörr. Globally optimal image partitioning by multicuts. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 31–44. Springer, 2011. 3, 6
- [11] Jonathan GA Lageard. *Dendrochronology*, pages 180–197. Springer, 2016. 1
- [12] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2
- [13] Zichen Liu, Jun Hao Liew, Xiangyu Chen, and Jiashi Feng. Dance: A deep attentive contour model for efficient instance segmentation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 345–354, 2021. 3
- [14] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv: Learning*, 2017. 6
- [15] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 6
- [16] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. *2016 Fourth International Conference on 3D Vision (3DV)*, pages 565–571, 2016. 3
- [17] Sida Peng, Wen Jiang, Huaijin Pi, Xiuli Li, Hujun Bao, and Xiaowei Zhou. Deep snake for real-time instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8533–8542, 2020. 3, 4, 6
- [18] Miroslav Poláček, Alexis Arizpe, Patrick Hüther, Lisa Weidlich, Sonja Steindl, and Kelly Swarts. Automation of tree-ring detection and measurements using deep learning. *bioRxiv*, 2022. 3
- [19] Giulia Resente, Alexander Gillert, Mario Trouillier, Alba Anadon-Rosell, Richard L Peters, Georg von Arx, Uwe von Lukas, and Martin Wilmking. Mask, train, repeat! artificial intelligence for quantitative wood anatomy. *Frontiers in plant science*, page 2526, 2021. 3
- [20] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing. 3
- [21] Georg von Arx and Marco Carrer. Roxas – a new tool to build centuries-long tracheid-lumen chronologies in conifers. *Dendrochronologia*, 32:290–293, 2014. 3
- [22] Georg von Arx and Hansjoerg Dietz. Automated image analysis of annual rings in the roots of perennial forbs. *International Journal of Plant Sciences*, 166:723 – 732, 2005. 3
- [23] Stef Weijers, Rob A. Broekman, and Jelte Rozema. Dendrochronology in the high arctic: July air temperatures reconstructed from annual shoot length growth of the circum-arctic dwarf shrub *cassiope tetragona*. *Quaternary Science Reviews*, 29:3831–3842, 2010. 1
- [24] Martin Wilmking, Allan Buras, Jiri Lehejcek, Jelena Lange, Rohan Shetti, and Ernst van der Maaten. Influence of larval outbreaks on the climate reconstruction potential of an arctic shrub. *Dendrochronologia*, 2018. 1
- [25] Adrian Wolny, Lorenzo Cerrone, Athul Vijayan, Rachele Tofanelli, Amaya Vilches Barro, Marion Louveaux, Christian Wenzl, Sören Strauss, David Wilson-Sánchez, Rena Lymbouridou, et al. Accurate and versatile 3d segmentation of plant tissues at cellular resolution. *Elife*, 9:e57613, 2020. 6, 7
- [26] Enze Xie, Pei Sun, Xiaoge Song, Wenhai Wang, Xuebo Liu, Ding Liang, Chunhua Shen, and Ping Luo. Polarmask: Single shot instance segmentation with polar representation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12190–12199, 2020. 3
- [27] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *ArXiv*, abs/1904.07850, 2019. 7