Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Live and Dead Knot Segmentation in X-ray Computed Tomography: A Comparative Study of Modern Networks

Shuai Zou

shuai.zou@ltu.se

Luleå University of Technology

Linus Olofsson

Luleå University of Technology

Oualid Burström

Luleå University of Technology

Thomas Reichert

Forstliche Versuchs- und Forschungsanstalt Baden-Württemberg

Olof Broman

Luleå University of Technology

Geir Isak Vestøl

Norwegian University of Life Sciences

Mohammad Jaber Hossain

Norwegian University of Life Sciences

Magnus Fredriksson

Luleå University of Technology

Johannes A. J. Huber

Wood K Plus - Kompetenzzentrum Holz GmbH

Research Article

Keywords: Computer vision, Deep learning, Knot morphology, Non-destructive evaluation, Norway spruce

Posted Date: April 21st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9134377/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License. Read Full License

Additional Declarations: No competing interests reported.

Live and Dead Knot Segmentation in X-ray Computed Tomography: A Comparative Study of Modern Networks

Shuai Zou[1*], Linus Olofsson[1], Oualid Burström[1], Thomas Reichert[2], Olof Broman[1], Geir Isak Vestøl[4], Mohammad Jaber Hossain[4], Magnus Fredriksson[1], Johannes A. J. Huber[3,1]

[1]Wood Science and Engineering, Luleå University of Technology, Forskargatan 1, Skellefteå, 93177, Sweden.
[2]Department Forest Utilisation, Forest Research Institute Baden-Wuerttemberg, Wonnhaldestrasse 4, Freiburg, 79100, Germany.
[3]Wood K Plus - Kompetenzzentrum Holz GmbH, Klagenfurter Straße 87–89, St. Veit an der Glan, 9300, Austria.
[4]Norwegian University of Life Sciences, Høgskoleveien 12, Ås, 5003, Norway.

*Corresponding author(s). E-mail(s): shuai.zou@ltu.se;
Contributing authors: linus.olofsson@ltu.se; oualid.burstrom@ltu.se; thomas.reichert@forst.bwl.de; olof.broman@ltu.se; geri.vestol@nmbu.no; mohammad.jaber.hossain@nmbu.no; magnus.1.fredriksson@ltu.se; j.huber@wood-kplus.at;

Abstract

Reliable segmentation of live and dead knots in X-ray computed tomography images is crucial for automated wood grading and sawing yield optimisation in the timber industry. Segmentation in wet-state X-ray computed tomography images remains challenging due to low contrast, noise, and the small, fragmented nature of dead-knot volumes. Under a controlled setting with a knot prior derived from the ground-truth whole-knot mask, which can be assumed

to be available in industrial evaluations, we benchmark representative modern semantic segmentation networks using a unified preprocessing, training, and evaluation protocol. We compare both region-overlap performance and boundary-localisation behaviour. To ensure consistent geometric assessment, we applied a lightweight, anatomy-informed post-processing to all predictions to stabilise the dead-knot boundary, and we complement region-overlap metrics with pith-referenced boundary-localisation metrics that quantify dead-knot boundary error. SegResNet achieves the strongest overall performance, including a live-knot Dice of 0.89 and a dead-knot Dice of 0.66, and also yields the most stable dead-knot boundary localisation. Overall, under the idealised setting where the knot area is constrained by ground truth, the residual encoder–decoder network (SegResNet) is particularly effective for dead-knot delineation on wet-state X-ray computed tomography slices.

Keywords: Computer vision, Deep learning, Knot morphology, Non-destructive evaluation, Norway spruce

1 Introduction

1.1 Knot Detection

Knots strongly affect the structural performance and value of sawn timber. They interrupt fibre continuity, alter local density and stiffness, and may introduce cracks and voids, especially in their dead portions (Koman et al. 2013). In this context, a *dead knot* denotes knot tissue originating from a branch that has died; the knot is no longer physiologically integrated with the surrounding stem, and the interface to adjacent wood is often weak or partially detached, which increases the likelihood of voids and cracks and reduces load transfer. Reliable, high-resolution mapping of knot geometry and, crucially, separation of live and dead regions can therefore improve grading decisions, optimise sawing strategies, and reduce waste across the wood value chain.

X-ray computed tomography (CT) is deployed in industrial sawmills to non-destructively reconstruct internal features before cutting and to support automated sawing optimisation at scale (Kobayashi et al. 2019; Ursella 2021). However, the

accuracy of current systems in locating the live–dead knot boundary is still limited, particularly under industrial constraints.

From a geometric standpoint, full 3D reconstruction with volumetric segmentation offers the most complete representation of knot continuity and morphology within the log. The operational burden of such fully volumetric pipelines is substantial: production-line throughput amplifies memory, computation, and integration costs, and these costs also slow large-scale experimental iteration. This motivates slice-oriented designs that retain use of 3D structure while avoiding full volumetric inference. Longuetaud et al. (2012) proposed a slice-wise connected-components approach for automatic knot detection and measurement to reduce memory usage, and Krähenbühl et al. (2016) tracked the knot pith and segmented knots on CT slices tangential to growth rings by extracting tangential images from the 3D volume rather than performing full volumetric segmentation. Similar accuracy–efficiency trade-offs between 3D and 2D segmentation networks have been reported in other imaging domains (Avesta et al. 2023).

In industrial green-state (including water) CT, live–dead knot segmentation is further complicated by imaging physics and anatomy. High moisture content in sapwood can shift grey levels toward those of knot tissue, reducing contrast between knots and surrounding wood. The class distribution is highly imbalanced: most pixels represent clear wood, while the regions that govern grading decisions are often dead knot portions and are small and fragmented. Finally, knots follow an anatomically continuous trajectory within the stem; local misclassifications therefore accumulate into systematic geometric errors along the live–dead boundary.

1.2 Automated Dead Knot Detection

Live and dead knot analysis for wood non-destructive testing has developed along two principal technical lines: (i) visual inspection based on high-resolution RGB imagery,

and (ii) volumetric analysis using CT. RGB-based methods are mature for post-sawing surface inspection of boards and panels, where the goal is to detect defects and often to distinguish live and dead knots after sawing. Classical pipelines include rule-based colour vision with thresholding, morphology, and geometric rules (Hu et al. 2004), and neurofuzzy segmentation with hand-crafted texture features (Ruz et al. 2005). More recent deep-learning systems treat *live knot* and *dead knot* as explicit categories in end-to-end detection and segmentation frameworks, for example using fully convolutional neural networks (CNNs) (He et al. 2019) or YOLO-type detectors and segmenters on industrial surface datasets (Wang et al. 2024). Because surface imaging cannot access internal 3D knot structure in logs, these approaches complement rather than replace CT-based pre-sawing log assessment.

Early CT studies relied on parametrised models with virtual sawing (Oja 2000), evaluating live/dead-related statistics by comparing reconstructed boards against real boards. Subsequent efforts targeted high-speed industrial CT, and the dead knot boundary was instead estimated directly in the volume while tracking knot axes across slices (Johansson et al. 2013). This remains a strong baseline that learning-based methods should ultimately surpass. At the same time, industrial CT imposes constraints such as coarse axial sampling and intensity variability, which motivates benchmarking modern networks under more controlled, higher-resolution conditions before confronting full industrial variability. Recent work has therefore shifted toward CT-based deep learning that learns texture–density–morphology features end-to-end (Giovannini et al. 2019), with the aim of improving robustness and enabling live–dead segmentation and boundary localisation.

In CT-based deep learning for live/dead knot analysis, Giovannini et al. (2019) proposed a two-stage CNN system: a whole-log semantic segmentation produces a knot probability volume and, using a pith prior combined with a Hough-based search, tracks individual knot axes; along each axis, a 3D region of interest (ROI) is resampled

into an oriented r–t–z (radial–tangential–longitudinal) volume on which live–dead discrimination (dead-knot border localisation) and diameter estimation are performed. The system was deployed on an industrial CT scanner at realistic anisotropic resolution and virtual sawing capability. Training targets for border and diameter estimation are obtained by sawing logs into boards, measuring knot attributes on board surfaces, and mapping these measurements back into the CT volume, aligning log-level predictions with physically relevant board measurements. However, such pipelines are usually evaluated as integrated systems. End-to-end performance conflates the intrinsic capacity of the segmentation model to separate live from dead tissue with upstream errors from knot detection, axis tracking, and knot contour extraction. A controlled benchmark that isolates pixel-level live–dead separation is therefore still missing.

1.3 Objective

Deep learning offers a direct route to optimise live–dead delineation from CT intensities by capturing local texture and context across multiple scales, which can mitigate intensity overlap and small object imbalance (Alzubaidi et al. 2021). However, the literature is dominated by defect vs. non defect (Gao et al. 2021) or whole-knot segmentation (Nordmark 2002; Krähenbühl et al. 2014). Pixel-level live/dead separation and measurement-oriented evaluation of boundary accuracy remain limited, and publicly available CT datasets with pixel-level annotations that distinguish live from dead knots are lacking. Existing CT-based wood imaging work tends to focus on species identification (Kobayashi et al. 2019), density estimation (Mahatara et al. 2024), or coarse defect labels, while differences in resolution, reconstruction pipelines, and annotation protocols hinder comparable baselines. We found no openly documented CT studies reporting pixel-level live/dead separation under wet-state conditions.

In this study we quantify how well modern semantic segmentation networks can separate live from dead knot tissue on 2D wet-state X-ray CT slices when knot

location is known. We therefore evaluate representative state-of-the-art networks in a ground-truth anatomy-guided setting where the model is only asked to classify pixels within given knot masks (priors), derived from the ground-truth whole-knot mask. This constrains the task to the live–dead interface itself and yields a reproducible best-case benchmark for boundary delineation. Using a unified preprocessing, training, and evaluation protocol, we compare both region-overlap accuracy and boundary-localisation behaviour, applying the same lightweight anatomy-informed post-processing to every prediction to stabilise the dead-knot boundary and reporting pith-referenced boundary-localisation errors alongside standard overlap metrics.

2 Materials and methods

2.1 Pipeline Overview

This work benchmarks live–dead knot segmentation on wet-state CT *slices* under a controlled setting in which knot localisation is treated as known. Figure 1 summarises the workflow. Given a wet-state slice and its ground-truth whole-knot mask $\mathcal{K}_{\mathrm{GT}}$, a *knot prior* $\mathcal{K}_{\mathrm{GT}}^{+}$ is constructed by dilating $\mathcal{K}_{\mathrm{GT}}$ so that it includes the knot itself and a surrounding neighbourhood. The CT slice and this prior are stacked as a two-channel input to a 2D semantic segmentation network that predicts three classes (non-knot, live knot, dead knot). All models were trained and evaluated with the same preprocessing, augmentation, loss, and split strategy.

To stabilise geometric assessment of the live–dead interface, a non-learned anatomy-informed post-processing is applied on every predicted mask, enforcing a single coherent live-to-dead transition along the knot axis (pith-to-knot direction), and suppressing spurious live–dead–live alternations. Performance is reported based on per-class overlap metrics and on a pith-referenced dead-knot boundary error.

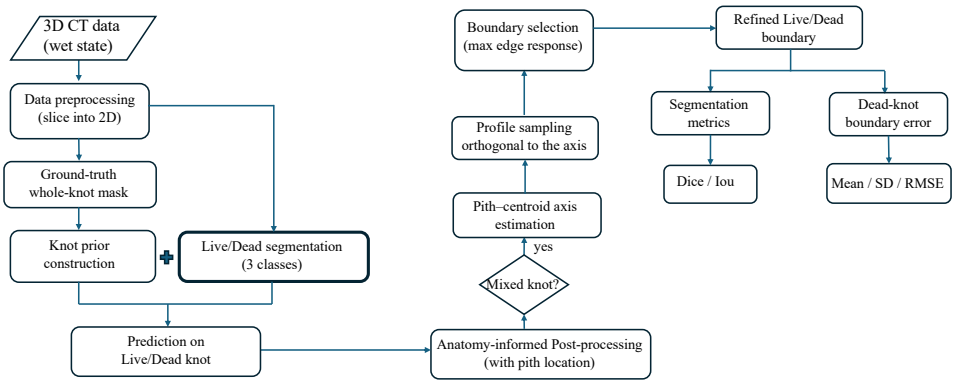


Fig. 1: Overview of the anatomy-guided benchmarking workflow for live/dead knot segmentation with post-processing: from wet-state 3D CT volumes, axial 2D slices were extracted. A knot prior was derived from the ground-truth whole-knot mask and stacked with the CT slice as a two-channel input for three-class (non-knot/live/dead) segmentation. The box with a thick border represents the learned segmentation models. Post-processing refined the live–dead boundary, and evaluation considered overlap metrics and dead-knot boundary error.

2.2 Specimens and CT Acquisition

The dataset comprised Norway spruce (*Picea abies* (L.) Karst.) from ten trees sampled across eight forest stands spanning northern to southern Sweden. Each tree was bucked into approximately 4 m logs and scanned in green state at Luleå University of Technology using a Microtec Mito industrial sawmill CT system (Figure 2) adapted for research. From each stem, whorl discs were sawn at fixed relative heights: 5% intervals from the butt up to 20% stem height and 10% intervals thereafter, with additional discs around 30% and 50% stem height to capture successive whorls. Each disc was CT-scanned in green state, then dried under controlled conditions to fibre saturation point to avoid cracking, and CT-scanned again in the dry state.

A helical cone-beam acquisition at 180 kV tube voltage and 5.55 mA current were used and tomographic reconstructions were isotropic with voxel size 0.5 mm, implementing the algorithm by (Katsevich 2004). The resulting 3D images with 16-bit

unsigned integers were computed using a machine-specific calibration to density in kg/m^3.

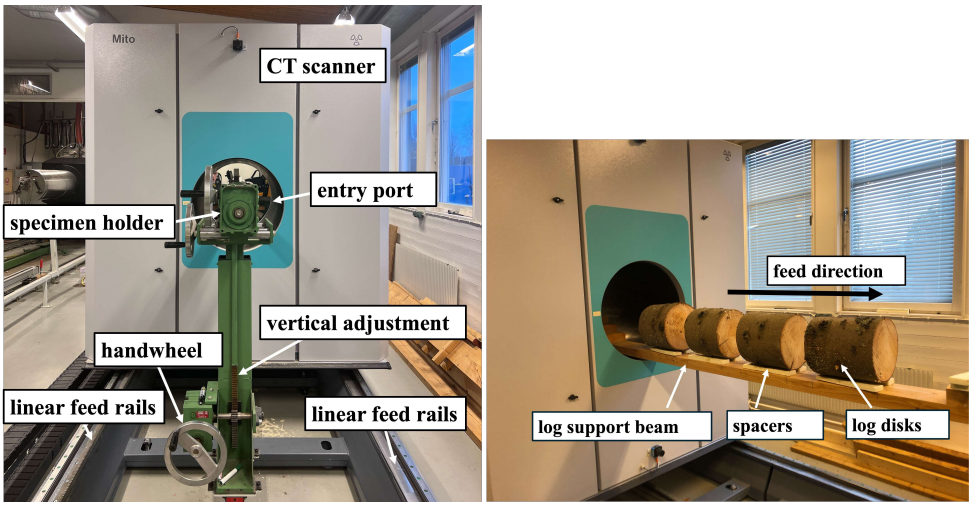


Fig. 2: Microtec Mito CT system: (a) Front view showing CT scanner with specimen holder and linear feed rails; (b) whorl-disc scanning setup with indicated feed direction, log discs, spacers and log support beam.

2.3 Annotation and Label Transfer

Ground truth was derived from physical measurements on each knot. In the wet state, knots on each disc were documented by azimuth (relative to a north reference), knot status at the mantle surface, and diameter. After drying, each knot was split radially through its pith to identify the dead-knot border and to measure the knot diameter at that border (Figure 3). These measurements defined five reference points per dry knot: the knot start at the pith, two points on the dead-knot border, and two points at the knot end at the cambium. The points were used in a custom annotation application to guide expert semi-automatic voxel labelling on dry-state CT images into three classes: non-knot, live knot, and dead knot.

To apply these labels to the wet state, dry volumes were rigidly aligned to their corresponding wet-state volumes using custom image registration software. The rigid registration estimated a rotation $R \in \mathrm{SO}(3)$ and translation $\mathbf{t} \in \mathbb{R}^3$,

$$T(\mathbf{x}) = R\mathbf{x} + \mathbf{t}, \tag{2.1}$$

mapping wet-grid coordinates $\mathbf{x} \in \Omega_{\mathrm{wet}}$ into dry-space. The registered dry image on the wet grid was

$$I_{\mathrm{dry}}^{\mathrm{reg}}(\mathbf{x}) = I_{\mathrm{dry}}\big(T(\mathbf{x})\big). \qquad \mathbf{x} \in \Omega_{\mathrm{wet}}, \tag{2.2}$$

Annotation volumes created on $I_{\mathrm{dry}}^{\mathrm{reg}}$ were stored on Ω_{wet} yielding co-registered 3D voxel-wise labels for all wet-state scans as

$$L : \Omega_{\mathrm{wet}} \to \{0, 1, 2\}, \tag{2.3}$$

where 0 = non-knot, 1 = live knot, 2 = dead knot.

We trained, validated, and tested only on wet-state scans because this matches the intended deployment setting and because wet-state images exhibited lower knot–sapwood contrast and stronger artefacts than dry-state images. The class distribution was strongly skewed: across trees, the no-knot class dominated with 98.4–99.2% of pixels, whereas live-knot and dead-knot pixels comprised only 0.73–1.39% and 0.09–0.48%, respectively.

2.4 Data Split and Slice Extraction

To prevent leakage, we split data 60/20/20 (train/validation/test) at the tree level. Each tree contributed 16 green (wet) 3D volumes, yielding 96 training volumes from six trees, 32 validation volumes from two trees, and 32 test volumes from two trees.

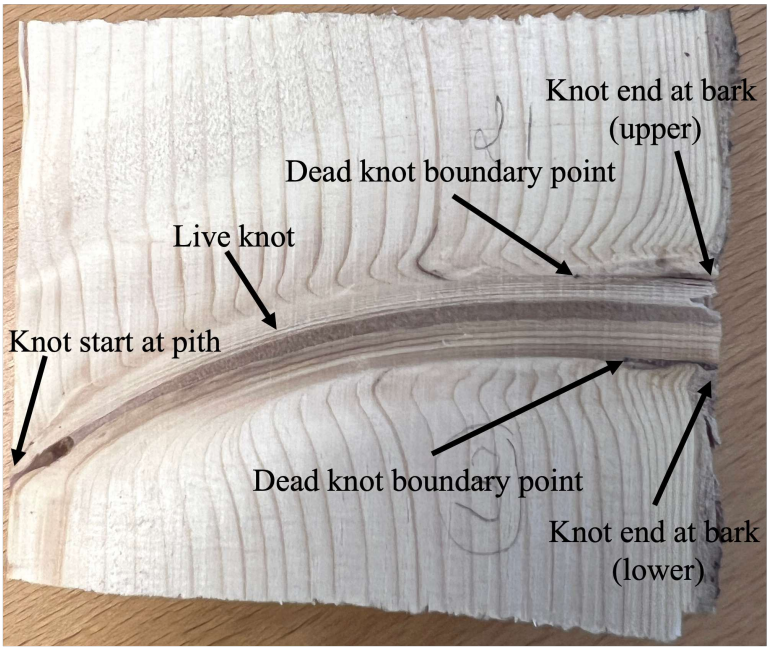


Fig. 3: Example of a disc radially cut through the knot pith. The five labelled points indicate the physical measurements used for annotation: the knot start at the pith, two dead-knot border points, and two knot-end points at the bark (upper and lower). At the far left, the pith of the log splitting into the pith of the knot is visible as a black line. Small inaccuracies during cutting could usually be corrected by inspecting the opposing piece. In the centre, the live knot is visible, with fibres (annual rings) clearly continuing into the knot. On the right, the dead knot is visible, with fibres terminating at the knot border and deviating to grow around the knot (out of the image plane). The dead-knot border was defined as the location where the first fibres began to grow around the knot. For difficult cases, this step was performed using a scalpel and a loupe (10× magnification), yielding an estimated dead-knot border accuracy of ±1 year. Because the border often differed above and below the knot, both positions were marked to provide a 3D description of the dead-knot border.

From each wet scan and its 3D label volume (I_{wet}, L) we extracted evenly spaced axial slices (2D image–label pairs) spanning the knot whorl region, discarding disc-end slices that contained no labelled knots. All CT volumes were reconstructed with an isotropic voxel size of 0.5 mm, so the pixel spacing was identical across scans.

The resulting datasets contained 15 920 training images, 5982 validation images, and 5760 test images.

2.5 Augmentation and Sampling

To mitigate low signal-to-noise ratio and contrast drift in wet-state CT, to improve the invariance to in-plane positioning of knots, we applied label-safe augmentations during training: geometric transforms were applied jointly to image and mask (nearest-neighbour resampling for masks), and intensity perturbations were applied to images only. Geometric augmentations were limited to flips, 90° rotations, and mild zoom with a scaling factor in [0.9, 1.1] to avoid unrealistic distortion of growth-ring geometry. Intensity jitter (brightness/contrast) was used to emulate acquisition variability. Random augmentations were applied during training; validation used deterministic preprocessing only.

To mitigate class imbalance, foreground-aware loss weighting was used: background-only slices were still sampled and forwarded to preserve stochasticity, but were assigned zero weight in the batch loss (i.e., only slices containing knot pixels contributed to the averaged loss) and therefore produced no gradients during backpropagation.

2.6 Segmentation Models

We benchmarked five 2D segmentation networks under the same training protocol. U-Net served as the reference encoder–decoder baseline (Figure 4) (Ronneberger et al. 2015). We then selected networks that probe distinct mechanisms relevant to wet-state live/dead separation: nested skip fusion (U-Net++, Zhou et al. (2018)), residual encoder–decoder optimisation (SegResNet, Myronenko (2018), MONAI implementation, Cardoso et al. (2022)), attention gating (Attention U-Net, Oktay et al. (2018)), and a CNN–Transformer hybrid with windowed self-attention (SwinUNet, Hatamizadeh et al. (2022)). Table 1 summarises the targeted challenges and expected trade-offs.

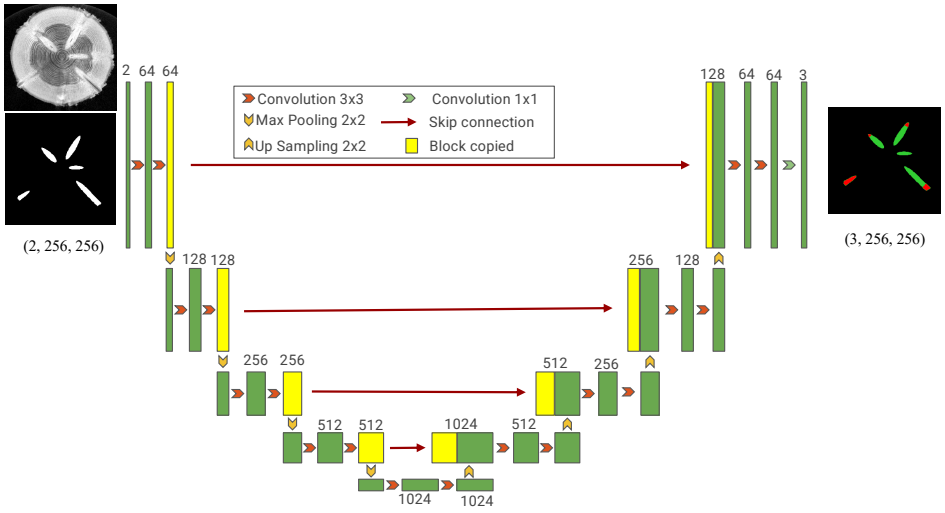


Fig. 4: U-Net architecture (adapted from Ronneberger et al. (2015)) used for the dead-knot segmentation task.

2.7 Knot Prior

A mask detecting the outline of the whole knot can be assumed to be available, even in industrial settings. We used a prior extending outwards from $\mathcal{K}_{\mathrm{GT}}$ as an *a priori* localisation cue because the most discriminative evidence for live/dead status lies near the interface: in live knots, fibres and growth rings continue into the knot, whereas in dead knots fibres terminate at the dead-knot border and adjacent fibres deviate to grow around the knot (Figure 3). The prior $\mathcal{K}_{\mathrm{GT}}^{+}$ was constructed by dilating $\mathcal{K}_{\mathrm{GT}}$ by 1 mm, which was based on validation experiments as a compromise between providing sufficient clearwood context at the boundary and avoiding reintroduction of large background regions and extreme class imbalance. Figure 5 illustrates the construction.

2.8 Segmentation Formulation

Given a preprocessed slice $I \in \mathbb{R}^{H \times W}$ extracted from I_{wet} (after export to $H = W = 256$), and binary prior $\mathcal{K}_{\mathrm{GT}}^{+}$, we form a two-channel input $u = \left[I_{\mathrm{norm}}, \mathcal{K}_{\mathrm{GT}}^{+}\right] \in$

Table 1: Representative networks and the architectural hypotheses probed.

Model	Core inductive bias	Targeted challenge	Main trade-off
U-Net Ronneberger et al. (2015)	Skip-connected encoder–decoder	Reference behaviour	Vulnerable to rare or fragmented dead-knot regions
SegResNet Myronenko (2018)	Residual blocks	Stable optimisation under low contrast	Risk of smoothing small fragments
SwinUNet Hatamizadeh et al. (2022)	Windowed self-attention	Context for oblique/discontinuous boundaries	Higher compute and tuning sensitivity
Attention U-Net Oktay et al. (2018)	Decoder attention gates	Foreground selectivity under sparsity	Marginal gain may be limited by strong prior
U-Net++ Zhou et al. (2018)	Nested skip fusion	Retention of thin/elongated structures	Higher architectural complexity

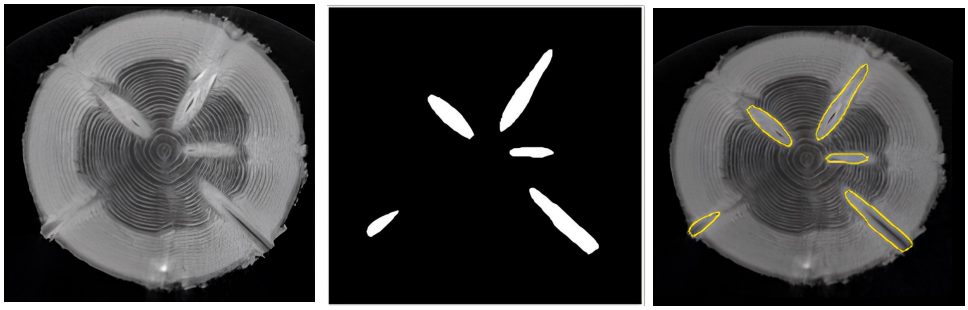


Fig. 5: Example of prior construction on a single slice: (a) input wet-state CT slice; (b) ground-truth whole-knot mask $\mathcal{K}_{\mathrm{GT}}$; (c) prior mask $\mathcal{K}^{+}_{\mathrm{GT}}$, visualised by its contour overlaid on the input slice.

$\mathbb{R}^{2\times H\times W}$, where I_{norm} denotes the intensity-normalised image. The network predicts per-pixel class probabilities

$$f_{\boldsymbol{\theta}}:\ \mathbb{R}^{2\times H\times W}\to\mathbb{R}^{3\times H\times W},\qquad \mathbf{p}=\mathrm{softmax}\big(f_{\boldsymbol{\theta}}(u)\big)\in[0,1]^{3\times H\times W}, \tag{2.4}$$

where $f_{\boldsymbol{\theta}}$ yields per-pixel class logits for mutually exclusive classes $\mathcal{C} = \{0, 1, 2\}$ (non-knot, live knot, dead knot), and softmax is applied across classes such that $\sum_{c \in \mathcal{C}} p_c(\mathbf{x}) = 1$ for every pixel $\mathbf{x}$. Predictions are

$$\hat{y}(\mathbf{x}) = \arg\max_{c \in \mathcal{C}} p_c(\mathbf{x}). \tag{2.5}$$

The prior is provided as an input channel rather than a hard mask so the network can still use surrounding context while being steered toward the boundary neighbourhood. The predicted map $\hat{y}$ was produced on the full field-of-view.

2.9 Training Protocol

All networks were trained with the same optimisation setup. Hyperparameters were selected via a sensitivity study on SegResNet and then fixed across networks to keep the benchmark reproducible. SegResNet was chosen because it exhibited instabilities that were representative of the overall training challenges in this task, making it a sensitive probe for identifying a conservative and robust operating regime. Hyperparameters identified in this way were then kept fixed across architectures; therefore, the comparison should be interpreted as performance under a unified conservative configuration rather than per-model best-case performance. The final configuration used learning rate 1×10^{-4}, weight decay 1×10^{-5}, no dropout, and an initial channel width of 32 feature channels.

We minimised a multi-class soft Dice loss computed on softmax probabilities, following Dice-based objectives for imbalanced segmentation introduced by Milletari et al. (2016).

$$\mathcal{L}_{\text{Dice}} = 1 - \frac{1}{3} \sum_{c=0}^{2} \frac{2 \sum_i p_{c,i}\, y_{c,i}}{\sum_i p_{c,i} + \sum_i y_{c,i} + \epsilon}, \tag{2.6}$$

Here, $p \in [0,1]^{3\times H\times W}$ denotes the per-slice network softmax output, and $y \in \{0,1,2\}^{H\times W}$ is the corresponding 2D ground-truth label map extracted from the 3D label volume L. $y_{c,i} \in \{0,1\}$ is one-hot encoding for each class c and pixel index i, and ϵ is added for numerical stability to avoid zero division when a class is absent or when p_c is nearly zero. This loss mitigated class imbalances. Optimisation used AdamW without learning-rate scheduling (Loshchilov and Hutter 2017). We trained with a batch size of 32 for both training and validation for 120 epochs, and report results using the checkpoint with the lowest validation loss. Training was performed on a workstation with an NVIDIA RTX A6000 GPU (48 GB), running Ubuntu 22.04.5. We used PyTorch 2.7.0 (CUDA 12.6) and MONAI 1.4.0 with 4 data-loader workers and a fixed random seed (0) for reproducibility.

2.10 Anatomy-informed Post-processing

Predictions before post-processing exhibited local label jitter, including repeated live–dead–live alternations within a single knot component. We therefore applied a network-agnostic geometric routine that calculated the dead-knot border with reference to the pith and enforced a single live-to-dead transition along the pith-to-knot direction.

Inputs were the predicted label map $\hat{y}$, the pith location $\mathbf{m}$ (manually annotated for each disc and reused for its slices), and connected components $\hat{\mathcal{K}}_j$ of the predicted knot mask $\hat{\mathcal{K}} = \hat{y} > 0$. For each $\hat{\mathcal{K}}_j$, if the minority knot class (dead/live) accounted for less than 10% of $\hat{\mathcal{K}}_j$ *and* had fewer than 120 pixels, we relabelled $\hat{\mathcal{K}}_j$ as the majority class. Otherwise $\hat{\mathcal{K}}_j$ was treated as mixed.

For a mixed $\hat{\mathcal{K}}_j$, we used the corresponding ground-truth component $\mathcal{K}_{j,\mathrm{GT}}$ for establishing a common ground for geometric evaluations. We calculated its centroid $\mathbf{g}$ and defined the unit axis

$$\mathbf{d} = \frac{\mathbf{g} - \mathbf{m}}{\|\mathbf{g} - \mathbf{m}\|}. \tag{2.7}$$

We estimated the dead-knot boundary position $\hat{s}_0$ by evaluating a set of candidate axial locations along the pith-to-knot axis. Specifically, we sampled equally spaced points $\mathbf{x}(t) = \mathbf{m} + t\,\mathbf{d}$ along the portion of the axis that intersects the predicted component $\hat{\mathcal{K}}_j$. At each candidate point, we extracted a short line segment through $\mathbf{x}(t)$ that is orthogonal to $\mathbf{d}$ and considered only the pixels of this segment that lie inside $\hat{\mathcal{K}}_j$. The predicted live/dead labels along this segment form a 1D profile; the dead-knot boundary corresponds to the location where this 1D label sequence changes from live to dead. To reduce one-pixel label flicker and tiny spurious islands, we smoothed this 1D label profile with a short 1D Gaussian filter (kernel size 3) and then computed its discrete gradient (finite differences) along the segment. We defined the transition *sharpness* at candidate t as the maximum absolute gradient magnitude of the smoothed profile and set $\hat{s}_0$ to the candidate axial location that maximised this sharpness.

If mixed knots contained live–dead–live alternations within a single knot component, all pixels in $\hat{\mathcal{K}}_j$ further away from $\hat{s}_0$, or from the perpendicular line through it, were relabelled to knot class *dead*, and all remaining pixels to class *live*, enforcing exactly one transition per mixed component (Figure 6).

2.11 Evaluation Metrics

We calculated per-class Dice and Intersection-over-Union (IoU) for live-knot and dead-knot pixels (Taha and Hanbury 2015; Everingham et al. 2010). Overall accuracy is omitted because background dominates the pixel count and would not reflect boundary quality. Dice and IoU take values in $[0, 1]$, where 1 indicates perfect agreement between prediction and ground truth.

For mixed knots, we calculated a pith-referenced dead-knot boundary error e_{s_0} in mm:

$$e_{s_0} = \hat{s}_0 - s_{0,\mathrm{GT}}, \tag{2.8}$$

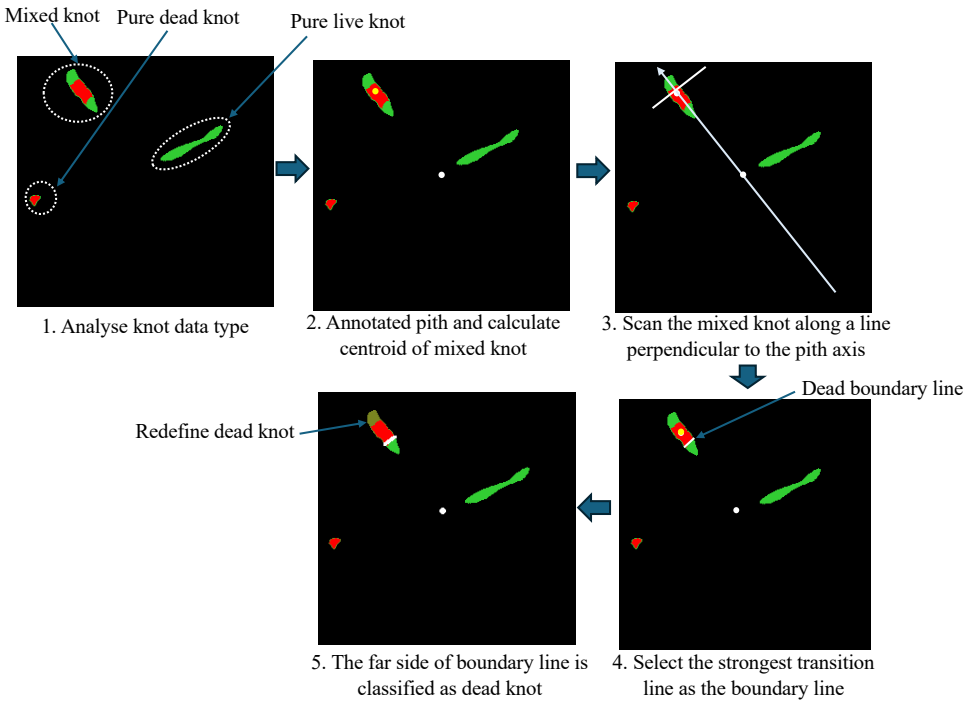


Fig. 6: Overview of post-processing. The example shows a knot whose prediction before post-processing exhibits a *live–dead–live* pattern along the primary axis (top row), which is converted by the post-processing into a geometrically consistent live–dead boundary (bottom row).

where $s_{0,\mathrm{GT}}$ is the true dead-knot boundary position, equivalently computed as $\hat{s}_0$ in each slice. Negative values indicate a shift towards the pith. We summarised e_{s_0} by its mean, standard deviation (SD), and root mean square error (RMSE) and calculated R^2 between $\hat{s}_0$ and $s_{0,\mathrm{GT}}$.

Each ground-truth connected knot component was categorised as pure-live, pure-dead, or mixed using the same criteria as in the post-processing step. Knot-type classification accuracy was then computed on these ground-truth components as the fraction whose predicted knot type matches the ground truth, reported separately for pure-live, pure-dead, and mixed. For knot-type classification, we report a confusion matrix with rows denoting the ground-truth knot type and columns denoting the

predicted knot type. For boundary localisation on mixed knots, we use an observed–predicted plot of the predicted axial boundary position $\hat{s}_0$ (y-axis) against the ground-truth position $s_{0,\mathrm{GT}}$ (x-axis). We also report histograms of the boundary error e_{s_0} to visualise its distribution beyond summary statistics.

3 Results and Discussion

3.1 Segmentation

Using the segmentation networks introduced in Section 2.6 and a fixed training protocol, we examined the predictions before post-processing to characterise differences in live/dead discrimination under wet-state CT. Table 2 summarises overlap metrics for the predictions before post-processing, and Figure 7 provides representative examples. Across all networks, live-knot overlap was higher than dead-knot overlap. Among the stronger networks, live-knot performance was comparable, whereas dead-knot segmentation remained more challenging. SegResNet achieved the highest Dice/IoU for both classes among the evaluated networks. This may reflect the improved optimisation stability of residual encoder–decoder designs under wet-state noise and low contrast. In contrast, SwinUNet and Attention U-Net achieved comparable live Dice but lower dead Dice than SegResNet. U-Net++ improved dead-knot overlap compared with U-Net while using fewer parameters. Its dead-knot Dice was nevertheless lower than SegResNet, which suggests that nested skip fusion alone may be insufficient to match the improvement obtained with the residual encoder–decoder under wet-state noise and low contrast. The qualitative examples in Figure 7 aligned with the quantitative trends: dead-knot regions were often under-segmented and appeared fragmented into isolated "islands", and in some cases small local label alternations were visible within slender components (see Figure 7, Ex. 5, U-Net++ and Ex. 4, SwinUNet). These artefacts motivated the anatomy-informed post-processing analysed next.

Table 2: Live/dead knot segmentation on wet-state slices from the test set (predictions before post-processing)

Model	Live Knot Dice	Live Knot IoU	Dead Knot Dice	Dead Knot IoU	Parameters ($\times 10^6$)
U-Net	0.5320	0.4143	0.3528	0.2602	6.4952
SegResNet	**0.8442**	**0.7659**	**0.5945**	**0.5172**	6.2994
SwinUNet	0.8302	0.7593	0.4497	0.3727	25.1570
Attention U-Net	0.8335	0.7563	0.4548	0.3757	7.9460
U-Net++	0.7424	0.6505	0.4833	0.3882	2.4102

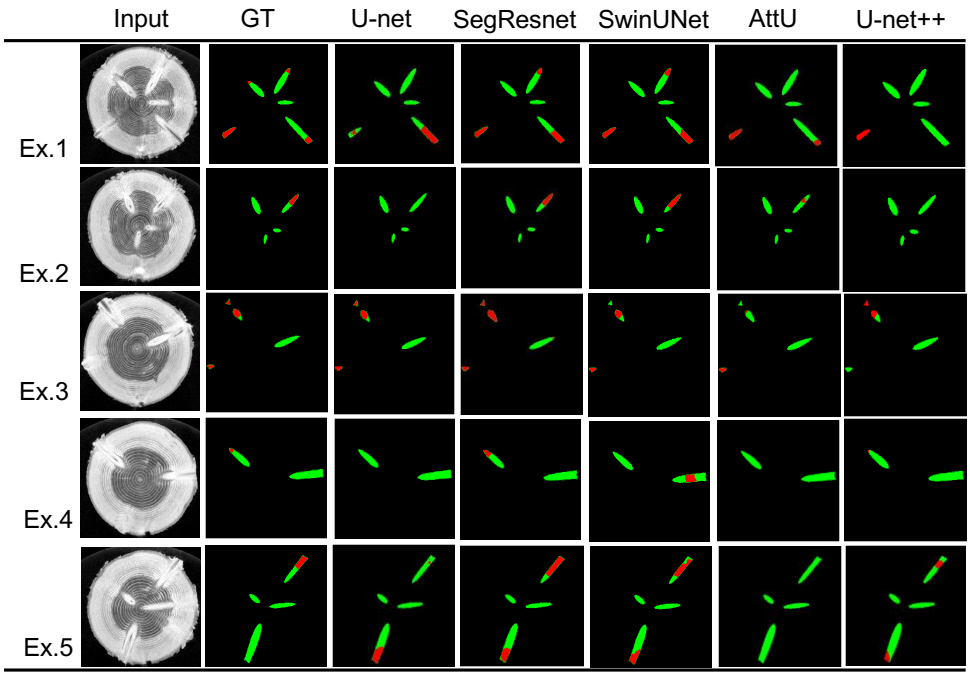


Fig. 7: Qualitative comparison of live/dead predictions before post-processing. Colour code: green = live knot; red = dead knot. Abbreviations: GT = ground truth; AttU = Attention U-Net.

3.2 Post-processing

Table 3 and Figure 8 show the updated overlap metrics and representative examples, respectively, after post-processing was applied. Post-processing reduces local live–dead label alternations and enforces a single live-to-dead transition per mixed component

by construction. Compared with the predictions before post-processing (Table 2), post-processing improved overlap metrics when averaged across the five networks, with a larger relative increase for dead knots (mean dead-knot Dice +31%) (Table 3). U-Net exhibits the largest improvement from predictions before post-processing to predictions after post-processing, which is consistent with its predictions before post-processing being more fragmented and showing local live–dead label alternations that are explicitly corrected by the proposed post-processing (see Figure 8, Ex. 1, U-Net).

As a result, the remaining errors were dominated by global failure modes that the geometric routine cannot correct, such as (i) missing dead-knot fragments in the predictions before post-processing (dead pixels absent altogether), (ii) spurious dead-knot predictions, where pixels were incorrectly labelled as dead (false positives for the dead-knot class), and (iii) incorrect inclusion or exclusion of knot pixels near the boundary that can shift the estimated transition location along the axis. Among the evaluated networks, SegResNet exhibited the highest performance after post-processing, indicating that it provided the most consistent predictions before post-processing for both knot completeness and boundary placement under wet-state CT.

Table 3: Live/dead knot segmentation on wet-state slices from the test set (predictions after post-processing)

Model	Live Knot Dice	Live Knot IoU	Dead Knot Dice	Dead Knot IoU
U-Net	0.8855	0.7945	0.6044	0.4331
SegResNet	**0.8928**	**0.8063**	**0.6647**	**0.5278**
SwinUNet	0.8845	0.7928	0.5741	0.4027
Attention U-Net	0.8923	0.8055	0.5964	0.4250
U-Net++	0.8736	0.7756	0.6206	0.4499

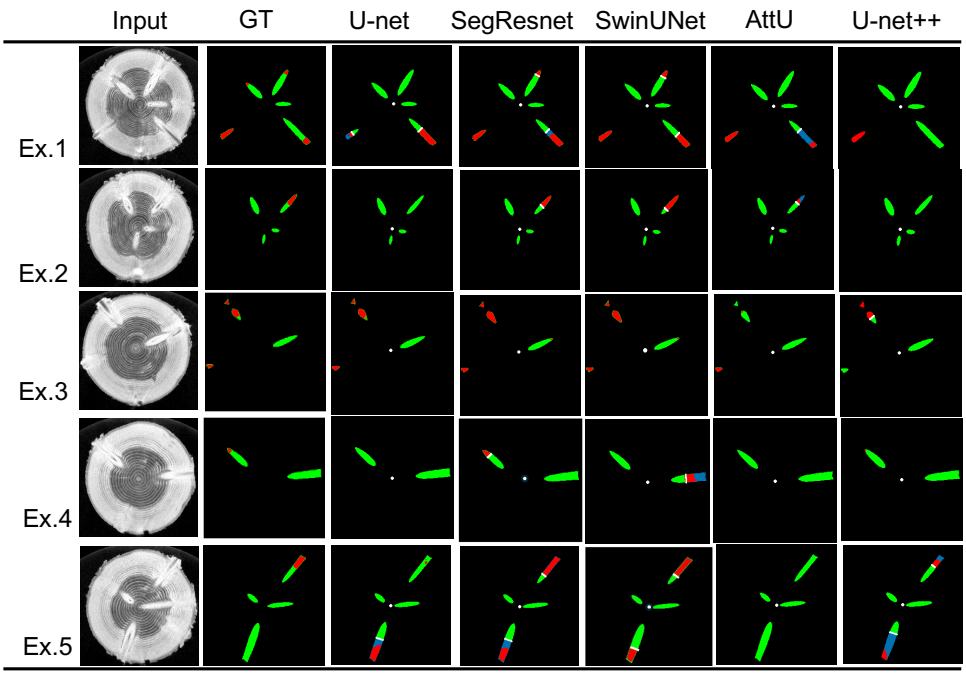


Fig. 8: Qualitative comparison of post-processed live/dead predictions. Colour code: green = live knot; red = dead knot; blue = dead knot relabelled by post-processing (from live to dead); white line = estimated live–dead boundary; white dot = pith. Abbreviations: GT = ground truth; AttU = Attention U-Net.

3.3 Boundary Localisation

Table 4 summarises knot-type classification accuracies and dead-knot boundary error (e_{s_0}) on mixed knots. For references, the values are compared to the metrics from Johansson et al. (2013), where the dead knot border error was reported for Scots pine (*Pinus sylvestris* L.) (n=65) under industrial log CT. For all networks, the mean e_{s_0} values were negative, implying an average pith-ward shift of the predicted boundary. This signed bias should, however, be interpreted with caution because the feasible error range is asymmetric and depends on the true boundary position: the dead-knot boundary cannot lie beyond the pith on the inner side or beyond the outer stem surface on the outer side. Consequently, larger $s_{0,\mathrm{GT}}$ values allow a wider negative error range but a narrower positive error range. The negative mean errors may therefore reflect

not only model behaviour but also these geometric constraints and the distribution of true $s_{0,\mathrm{GT}}$ values in the dataset. SegResNet yielded the most consistent boundary localisation (lowest SD/RMSE), whereas AttentionUNet showed the smallest mean bias but the largest dispersion, indicating less stable boundary placement across mixed cases.

Table 4: Classification accuracy computed on all pairs and summary statistics on dead-knot boundary error e_{s_0} on mixed knots only (computed from post-processed predictions on the test set)

Model	Accuracy			e_{s_0}			
	pure live (%)	**pure dead (%)**	**mixed (%)**	**Mean (mm)**	**SD (mm)**	**RMSE (mm)**	R^2
U-Net	90.63	55.92	53.27	−2.97	17.14	17.38	0.02
SegResNet	88.20	**76.75**	**68.80**	−5.65	**14.38**	**15.48**	0.14
SwinUNet	88.25	55.85	48.12	−4.85	17.69	18.33	0.10
Attention U-Net	92.29	52.91	50.51	**-0.97**	18.77	18.79	**0.17**
U-Net++	**93.23**	54.47	55.70	−1.53	16.66	16.72	0.02
Johansson et al. (2013)	–	–	–	−4.0	12	12	0.19

The knot-type classification results provided complementary context (Table 4). Knot-type classification accuracy was consistently highest for pure-live cases across all networks, whereas accuracy for pure-dead and mixed cases was lower and showed inter-model gaps. SegResNet achieved the highest accuracy on both pure-dead and mixed knots. Figure 9 provides diagnostic plots for SegResNet, including a knot-type confusion matrix and an observed–predicted plot of the axial boundary position s_0 on mixed knots. The observed–predicted plot in Figure 9 shows vertical banding, indicating that many samples share identical or near-identical $s_{0,\mathrm{GT}}$ values. This is expected given the discrete pixel grid of the CT slice and the boundary-localisation procedure, which searches over a discretely sampled set of candidate axial positions when computing both $s_{0,\mathrm{GT}}$ and $\hat{s}_0$. Beyond this discretisation, the plot shows that errors become increasingly negative for larger $s_{0,\mathrm{GT}}$, i.e., $\hat{s}_0$ more often falls below the identity line when the true boundary lies further outward. This trend is also supported

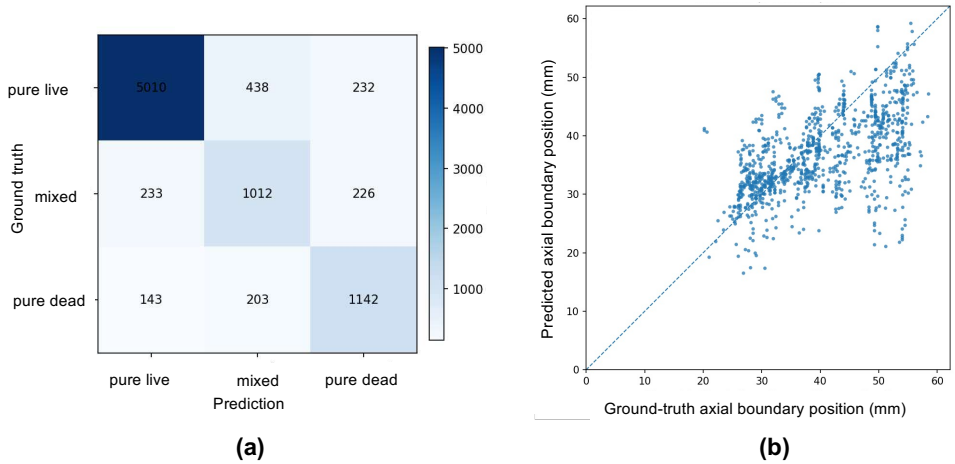


Fig. 9: Diagnostic plots for SegResNet, selected as the network with the lowest SD/RMSE of e_{s_0} on mixed knots (Table 4). (a) knot-type confusion matrix on the test set. (b) observed–predicted plot of the axial live–dead boundary position s_0 (mm) on mixed knots (GT=mixed and Pred=mixed), measured along the pith-to-knot direction $\mathbf{d}$ defined from the ground-truth component (Eq. 2.7) and reused for both ground truth and prediction; the dashed line indicates perfect agreement ($y = x$).

by the error histograms in Figure 10, where the distribution of e_{s_0} shifts progressively toward more negative values for larger $s_{0,\mathrm{GT}}$ and appears to become broader. Part of this pattern may arise from the asymmetric, position-dependent bounds of the signed error itself. The slightly positive average error in the small-$s_{0,\mathrm{GT}}$ group is consistent with this interpretation. Additional contributors may include weaker or more fragmented dead-knot cues farther from the pith, as well as the tendency of the slice-wise 2D formulation and the post-processing step to favour an earlier transition when the evidence is ambiguous. In practical terms, such bias would shorten the predicted sound-knot length while enlarging the predicted dead portion.

3.4 Performance Comparison

A strict numerical comparison of our results with previously published industrial log-CT pipelines is not straightforward, mainly because the task setup and the reported metrics are not fully aligned. In particular, our benchmark isolates pixel-level live/dead

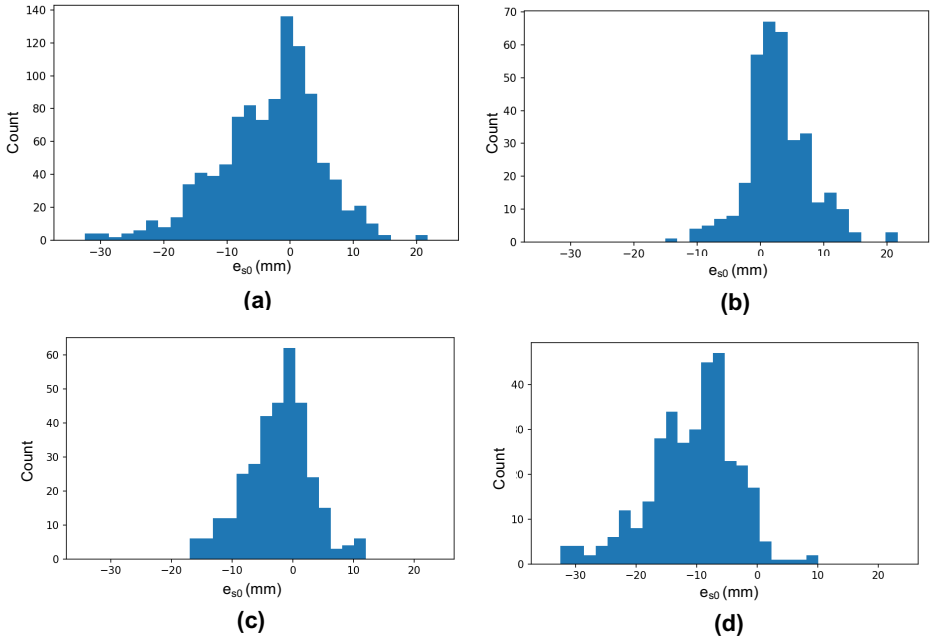


Fig. 10: Histograms of dead-knot boundary error $e_{s_0} = \hat{s}_0 - s_{0,\mathrm{GT}}$ on mixed knots. (a) All mixed–mixed samples. (b) Small $s_{0,\mathrm{GT}}$ ($\leq 34\,\mathrm{mm}$). (c) Medium $s_{0,\mathrm{GT}}$ (34–46 mm). (d) Large $s_{0,\mathrm{GT}}$ ($> 46\,\mathrm{mm}$). Samples in (b)–(d) are grouped by $s_{0,\mathrm{GT}}$ terciles (33%/66% quantiles).

discrimination on 2D slices under a controlled setting with a ground-truth-derived knot prior, whereas several earlier approaches operate on whole logs and leverage 3D context (including axial knot extent and continuity). Moreover, the target quantities and error measures reported in prior work are not always defined identically to those considered here. We therefore use prior work primarily as qualitative reference points.

Johansson et al. (2013) proposed an automatic knot-detection algorithm for high-speed industrial log CT. On the Scots pine data (our dataset is Norway spruce), they reported a dead-knot-border error with a mean bias of about $-4\,\mathrm{mm}$, a SD and RMSE of about 12 mm, and $R^2 \approx 0.19$ (sample size $n = 65$). Importantly, their dead-knot border is defined as a pith-referenced distance associated with the sound–dead transition and is operationalised via the axial position where the knot size reaches

its maximum, which is not strictly equivalent to our slice-wise pith-to-knot boundary position $\hat{s}_0$ and the corresponding error $e_{s_0} = \hat{s}_0 - s_{0,\mathrm{GT}}$. Nevertheless, their results provide a useful order-of-magnitude reference. In our wet-state disc benchmark, the best-performing network achieved an RMSE of 15.48 mm for e_{s_0} on mixed knots (Table 4). The remaining gap is consistent with the intrinsic limitation of our strictly slice-wise 2D formulation: without exploiting 3D knot continuity and axial context, local ambiguities and fragmented dead-knot cues can accumulate into larger deviations in boundary placement. This suggests that incorporating 3D information is necessary to achieve production-grade pipelines. To avoid the computationally demanding task of accounting for the full 3D geometry, windowed or other partially 3D approaches may be appropriate, e.g. 2.5D Bajić et al. (2025).

Giovannini et al. (2019) proposed a two-stage deep-learning pipeline for industrial log CT, combining knot-axis tracking with an oriented 3D ROI for sound/dead discrimination. On Scots pine, they reported sound/dead classification accuracies of 88.5% (sound) and 83.5% (dead). In our controlled wet-state disc benchmark, the corresponding knot-type accuracies for SegResNet were 88.20% on pure-live knots, 76.75% on pure-dead knots, and 68.80% on mixed knots (Table 4). These figures are not directly comparable due to differences in task definition and evaluation unit (binary sound/dead status evaluated at board-intersection points versus three-way knot-type classification on slice-wise connected components from pixel-level segmentations under wet-state CT). We therefore use the results of Giovannini et al. (2019) only as a coarse reference for classification performance.

4 Conclusion

We present a controlled benchmark for live–dead knot segmentation in wet-state X-ray CT slices with a ground-truth-derived knot prior provided, enabling fair comparison of modern 2D segmentation networks with a unified training and evaluation protocol.

Using both region-overlap and pith-referenced boundary-localisation metrics, with the same anatomy-informed post-processing applied to all networks, SegResNet achieved the highest overlap scores and the lowest boundary-localisation dispersion on mixed knots. Motivated by the practical efficiency of slice-wise 2D processing for long logs, our benchmark provides a reproducible baseline and a quantitative reference for future 2D refinements, 3D extensions, and industrial-scale methods.

Future work could explore integrating wood's biological and physical priors into deep learning models to improve the accuracy and robustness of live–dead knot segmentation. For example, prior knowledge such as growth-ring density distributions, fibre-orientation patterns, and species-specific characteristics may be incorporated as additional inputs or as structural constraints. Another natural extension is to move beyond the present slice-wise 2D formulation toward 3D volumetric architectures, either fully or in a windowed manner. A combination of wood-science priors, 3D reasoning, and realistic voxel sampling could potentially reduce the reliance on intensity and texture cues alone, enabling more accurate boundary identification, more biologically plausible segmentations, and closer alignment with industrial-grade automated wood inspection and grading.

Statements and Declarations

Funding

This study is part of the project *Wood Artificial Intelligence - Modelling of Knots with Computed Tomography (WAI-KnotCT)* and was kindly funded by Sweden's Innovation Agency (Vinnova) through the grant Bioeconomy in the North. The research was also partially funded by the Formas project 2022-00469 *Increasing the share of long-lifecycle wood products by extracting the value of data in the sawmill industry*, and by the Swedish Energy Agency, Energimyndigheten, Sweden and Formas, Sweden via the strategic innovation programme RE:Source. The authors also thank for the

support from the CT WOOD research programme in Skellefteå, Sweden, which is generously supported by the Swedish sawmilling industry, Skellefteå municipality and Luleå University of Technology, Sweden.

Competing Interests

The authors have no competing interests to declare that are relevant to the content of this article.

Acknowledgements

The authors thanks Lars Hansson and Olov Karlsson, Luleå University of Technology, for their work in preparing the dataset and guidence, respectively.

Author Contributions

Shuai Zou: conceptualisation, investigation, methodology, formal analysis, writing – original draft; **Linus Olofsson**: supervision, writing – review & editing, data curation; **Oualid Burström**: software; **Thomas Reichert**: data curation; **Olof Broman**: data curation; **Geir Isak Vestøl**: data curation; **Mohammad Jaber Hossain**: data curation; **Magnus Fredriksson**: supervision, writing – review & editing; **Johannes A. J. Huber**: conceptualisation, supervision, writing – review & editing.

Data Availability

The datasets generated and/or analysed during the current study are not publicly available due to proprietary constraints and agreements with project partners. Access may be granted by the corresponding author on reasonable request, subject to any applicable restrictions.

References

Alzubaidi L, Zhang J, Humaidi AJ, et al (2021) Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. Journal of Big Data 8(1):53. https://doi.org/10.1186/s40537-021-00444-8

Avesta A, Hossain S, Lin M, et al (2023) Comparing 3D, 2.5D, and 2D approaches to brain image auto-segmentation. Bioengineering 10(2):181. https://doi.org/10.3390/bioengineering10020181

Bajić B, Huber JAJ, Neyses B, et al (2025) Sparse view tomographic reconstruction of elongated objects using learned primal-dual networks. Engineering Applications of Artificial Intelligence 162:112295. https://doi.org/10.1016/j.engappai.2025.112295

Cardoso MJ, et al (2022) Monai: An open-source framework for deep learning in healthcare. https://doi.org/10.48550/arXiv.2211.02701, preprint at https://arxiv.org/abs/2211.02701

Everingham M, Van Gool L, Williams CKI, et al (2010) The Pascal Visual Object Classes (VOC) challenge. International Journal of Computer Vision 88(2):303–338. https://doi.org/10.1007/s11263-009-0275-4

Gao M, Song P, Wang F, et al (2021) A novel deep convolutional neural network based on resnet-18 and transfer learning for detection of wood knot defects. Journal of Sensors 2021:4428964. https://doi.org/10.1155/2021/4428964

Giovannini S, Boschetto D, Vicario E, et al (2019) Improving knot segmentation using deep learning techniques. Tech. Rep. FPL-GTR-272, USDA Forest Service, Forest Products Laboratory, URL https://www.fpl.fs.fed.us/documnts/fplgtr/fpl_gtr272.pdf, paper presented at the 21st International Nondestructive Testing and Evaluation of Wood Symposium, Freiburg, Germany, 24–27 September 2019

Hatamizadeh A, Nath V, Tang Y, et al (2022) Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. https://doi.org/10.48550/arXiv.2201.01266, preprint at https://arxiv.org/abs/2201.01266

He T, Liu Y, Xu C, et al (2019) A fully convolutional neural network for wood defect location and identification. IEEE Access 7:123453–123462. https://doi.org/10.1109/ACCESS.2019.2937461

Hu C, Tanaka C, Ohtani T (2004) Locating and identifying sound knots and dead knots on sugi by the rule-based color vision system. Journal of Wood Science 50(2):115–122. https://doi.org/10.1007/s10086-003-0549-3

Johansson E, Johansson D, Skog J, et al (2013) Automated knot detection for high speed computed tomography on pinus sylvestris l. and picea abies (l.) karst. using ellipse fitting in concentric surfaces. Computers and Electronics in Agriculture 96:238–245. https://doi.org/10.1016/j.compag.2013.06.003

Katsevich A (2004) An improved exact filtered backprojection algorithm for spiral computed tomography. Advances in Applied Mathematics 32(4):681–697. https://doi.org/10.1016/S0196-8858(03)00099-X

Kobayashi K, Hwang SW, Okochi T, et al (2019) Non-destructive method for wood identification using conventional x-ray computed tomography data. Journal of Cultural Heritage 38:88–93. https://doi.org/10.1016/j.culher.2019.02.001

Koman S, Feher S, Abraham J, et al (2013) Effect of knots on the bending strength and the modulus of elasticity of wood. Wood Research 58(4):617–626

Krähenbühl A, Kerautret B, Debled-Rennesson I, et al (2014) Knot segmentation in 3d ct images of wet wood. Pattern Recognition 47(12):3852–3869. https://doi.org/10.1016/j.patcog.2014.05.015

Krähenbühl A, Roussel JR, Kerautret B, et al (2016) Robust knot segmentation by knot pith tracking in 3d tangential images. In: Computer Vision and Graphics. Springer, pp 581–593, https://doi.org/10.1007/978-3-319-46418-3_52

Longuetaud F, Mothe F, Kerautret B, et al (2012) Automatic knot detection and measurements from x-ray CT images of wood: A review and validation of an improved algorithm on softwood samples. Computers and Electronics in Agriculture 85:77–89. https://doi.org/10.1016/j.compag.2012.03.013

Loshchilov I, Hutter F (2017) Decoupled weight decay regularization. https://doi.org/10.48550/arXiv.1711.05101, arXiv:1711.05101

Mahatara D, Campelo F, Houle L, et al (2024) Ctring: An r package to extract wood density profiles from computed tomography images of discs and logs. Dendrochronologia 88:126274. https://doi.org/10.1016/j.dendro.2024.126274

Milletari F, Navab N, Ahmadi SA (2016) V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 Fourth International Conference on 3D Vision (3DV), pp 565–571, https://doi.org/10.1109/3DV.2016.79

Myronenko A (2018) 3d mri brain tumor segmentation using autoencoder regularization. https://doi.org/10.48550/arXiv.1810.11654, preprint at https://arxiv.org/abs/1810.11654

Nordmark U (2002) Knot identification from ct images of young pinus sylvestris sawlogs using artificial neural networks. Scandinavian Journal of Forest Research 17(1):72–78. https://doi.org/10.1080/028275802317221109

Oja J (2000) Evaluation of knot parameters measured automatically in ct-images of norway spruce (picea abies (l.) karst.). Holz als Roh- und Werkstoff 58(5):375–379. https://doi.org/10.1007/s001070050448

Oktay O, Schlemper J, Folgoc LL, et al (2018) Attention U-net: Learning where to look for the pancreas. https://doi.org/10.48550/arXiv.1804.03999

Ronneberger O, Fischer P, Brox T (2015) U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, pp 234–241, https://doi.org/10.1007/978-3-319-24574-4_28

Ruz GA, Estevez PA, Perez CA (2005) A neurofuzzy color image segmentation method for wood surface defect detection. Forest Products Journal 55(4):52–58

Taha AA, Hanbury A (2015) Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. BMC Medical Imaging 15(1):29. https://doi.org/10.1186/s12880-015-0068-x

Ursella E (2021) In-line industrial computed tomography applications and developments. PhD thesis, Università Ca' Foscari Venezia

Wang R, Chen Y, Liang F, et al (2024) Bpn-yolo: A novel method for wood defect detection based on yolov7. Forests 15(7):1096. https://doi.org/10.3390/f15071096

Zhou Z, Rahman Siddiquee MM, Tajbakhsh N, et al (2018) Unet++: A nested u-net architecture for medical image segmentation. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support. Springer, p 3–11, https://doi.org/10.1007/978-3-030-00889-5_1