Clearcut Detection Using Sentinel-1 SAR Data

Qamcom Research and Technology

Abstract—**Reliable detection of forest clearcuts is a critical requirement for regulatory enforcement, sustainable forest management, and environmental oversight. In regions affected by persistent cloud cover, limited daylight, and seasonal snow, optical satellite imagery often fails to provide consistent and timely observations, creating gaps in operational monitoring. This white paper presents an operational clearcut detection solution based on synthetic aperture radar (SAR) data from the Sentinel-1 mission combined with a neural network specifically optimized for radar imagery. Sentinel-1 was selected to meet key system requirements for cost-effective and long-term reliable monitoring services through stable, open-access radar acquisitions, even though this was not the most optimal radar for this application. To our knowledge, this is among the first operational efforts to perform large-scale clearcut detection using Sentinel-1 SAR change pairs and a learning-based, radar-tailored segmentation model as the primary sensing modality. The proposed system leverages dual-polarized SAR data and a multi-scale feature pyramid architecture to identify forest structural changes independently of weather and illumination conditions. Extensive validation throughout Sweden demonstrates that the system detects the majority of clearcuts while maintaining a low rate of false alarms for operationally relevant harvest sizes. Overall, the results indicate a scalable, robust, and deployment-ready capability for continuous, national-scale forest monitoring, with straightforward adaptation to new regions through model fine-tuning.**

INTRODUCTION

Operational forest monitoring requires reliable, repeatable and timely data products to support enforcement, planning, and sustainability objectives. While optical satellite imagery has historically been the primary data source for detecting forest changes, its effectiveness is significantly reduced in northern and temperate regions due to persistent cloud cover, limited daylight during winter, and snow-covered terrain. These limitations create blind spots in monitoring workflows and increase reliance on delayed or incomplete information. Synthetic aperture radar (SAR) offers a practical alternative for operational use. Sentinel-1 provides consistent, large-area coverage with a 12-day revisit time, 10-meter spatial resolution, and dual polarization (VV and VH). Because SAR operates independently of weather and lighting conditions, it enables uninterrupted monitoring throughout the year. This capability is particularly valuable for forest agencies and commercial operators seeking continuous situational awareness.

While longer-wavelength SAR systems can provide more direct sensitivity to forest biomass, the Sentinel-1 mission was selected primarily for its long-term operational continuity and guaranteed data availability. In addition, its open-access data policy supports cost-effective monitoring services at national scale. Mission longevity and stability are essential requirements for building sustainable, production-grade monitoring services, even when this constrains the system to C-band radar data. C-band backscatter is more strongly influenced by branches, undergrowth, and environmental conditions, which increases the complexity of forest change detection. Nevertheless, by leveraging consistent Sentinel-1 acquisitions and advanced neural network models, the proposed system successfully extracts clearcut signals from these challenging observations. The achieved performance demonstrates that reliable, large-scale clearcut detection can be realized even under the physical limitations imposed by C-band SAR, provided that long-term, stable data streams and appropriate learning-based methods are employed. This is an example of designing and optimizing an end-to-end solution around operational system requirements, even when the sensing modality is not considered the optimal choice for the task.

Figure 1 provides an overview of the workflow. First, $50\,\text{km} \times 50\,\text{km}$ radar images are collected from regions across Sweden over multiple dates. Next, before–after image pairs are created by combining radar acquisitions from different times. These paired images are then used to train our proprietary neural networks, which produce a clearcut mask during inference. This end-to-end workflow is designed for operational use: automated data ingestion, consistent preprocessing, robust inference under all-weather conditions, and outputs that can be used directly for monitoring and reporting.

METHOD

The presented solution is designed as an end-to-end, operational clearcut detection pipeline using multitemporal Sentinel-1 SAR data.

Data Acquisition

The first step involves acquiring Sentinel-1 data. We retrieve Sentinel-1 radar imagery through a data acquisition platform that interfaces with the Copernicus mission developed by Spacemetric. The system is configured using parameters such as region coordinates, area of interest, and time range, enabling seamless and automated data collection. The acquired imagery is then preprocessed to ensure it is consistent and analysis-ready. This preprocessing includes radiometric calibration, geometric alignment, and normalization to reduce variability across acquisitions and environmental conditions. Signal processing and geometric calibration further ensure that the data is accurately aligned with the ground geometry, enabling a reliable mapping between ground objects and image pixels. The radar signal values represent the normalized backscatter coefficient, denoted by σ. We use both VV and VH polarization channels at a *10m resolution per pixel* to capture complementary structural information from the forest canopy and the

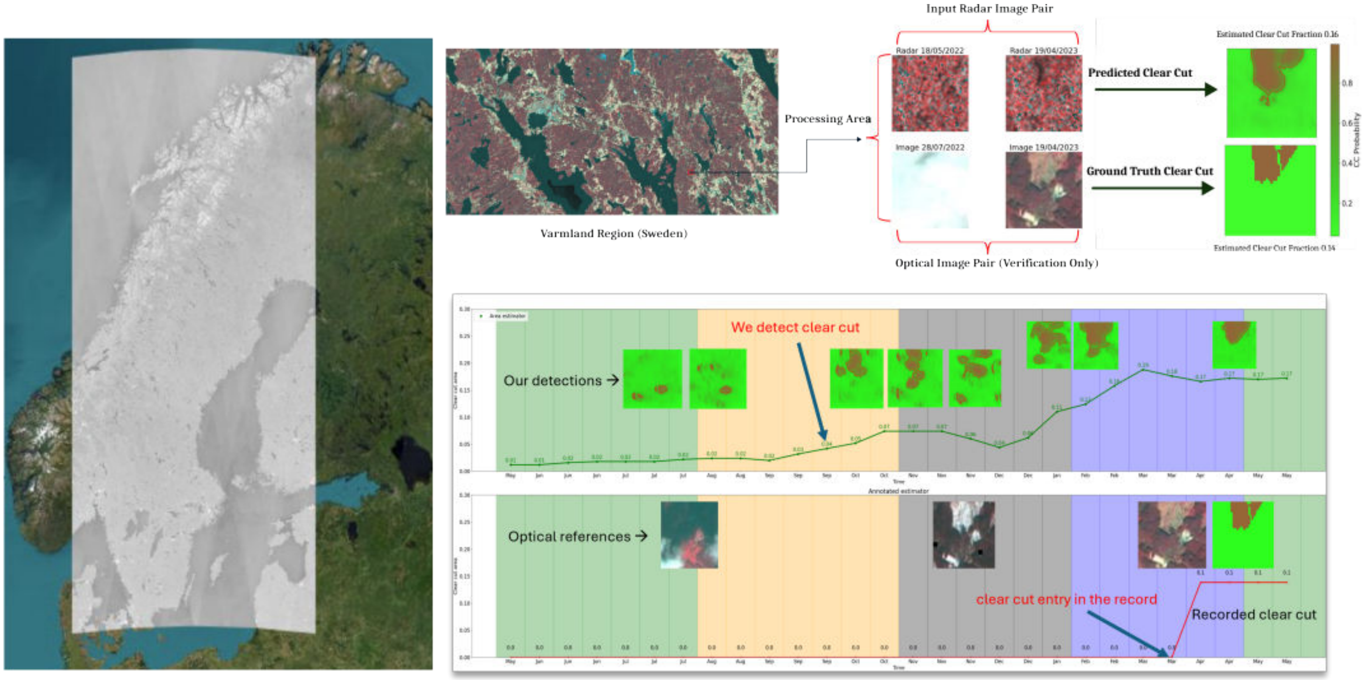

Fig. 1. Clearcut Detection Workflow

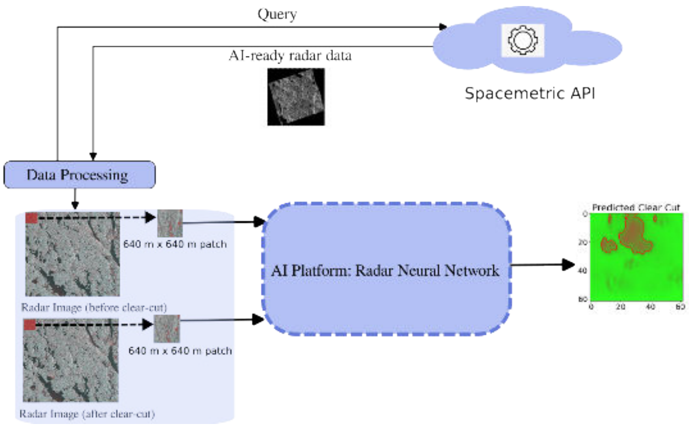

Fig. 2. Multi-scale feature pyramid network used for clearcut segmentation from before–after Sentinel-1 patches.

ground layer. Finally, we operate a downstream receiver that connects to the Spacemetric endpoint via a REST API to fetch these AI-ready radar images, which are subsequently used to train our model.

The other component of model training is obtaining ground-truth masks. These are provided by *Skogsstyrelsen (Swedish Forest Agency)*, which records clearcuts occurring in Sweden and maintains them in a georeferenced map. In this map, each pixel encodes the date on which the clearcut occurred.

Data Preprocessing

We formulate clearcut detection as a change detection problem. The goal is to identify changes in the forest canopy between two acquisition dates; such canopy changes serve as a proxy for whether a clearcut occurred in the intervening period. To do this, we use the upstream Spacemetric system to retrieve a time series of satellite images for a given forest region. From this time series, we construct image pairs (I_{t_1} and I_{t_2}) with $t_1 < t_2$, motivated by the idea that any clearcut activity would manifest as a detectable change between the earlier and later observations. We refer I_{t_1} as radar image before clearcut and I_{t_2} as the one after clearcut. For supervision, we use the georeferenced clearcut masks provided by Skogsstyrelsen. For each image pair, we spatially crop the Skogsstyrelsen mask to the same region and temporally filter it to the same interval $[t_1, t_2]$. This yields a label mask indicating where clearcut events are recorded as having occurred between the two dates. As a result, each training example consists of:

- an image pair: before and after observations of the same area, and,
- a corresponding ground-truth mask derived from Skogsstyrelsen for the interval between those dates.

Model Training

To enable an optimized training process under computational constraints, it is suboptimal to process full-sized radar scenes end-to-end. Instead, we adopt a *patch-based* training strategy. Concretely, we use *64×64* image patches, corresponding to approximately $640m^2$ of forest area per patch (as shown in Figure 2). For each *before clearcut* and *after clearcut* image pair, we sample a large number of co-registered patch pairs from the same geographic locations. These paired patches, together with their corresponding label patches, form the training dataset used to learn the model parameters. A multi-scale feature pyramid neural network

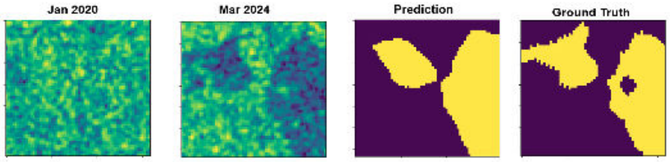

Fig. 3. Generated evaluation mask from before and after cut radar image validated with ground truth mask.

(Figure 2) that is built leveraging the Unet architecture forms the core of the detection system. This architecture is well suited for operational forestry use because clearcuts occur across a wide range of spatial scales, from small openings to large contiguous harvest areas. By learning features at multiple resolutions, the feature pyramid design captures both broad contextual information (useful for larger clearcuts) and fine spatial detail (needed to delineate boundaries precisely). Given a *before* and *after* patch as input, the network outputs a pixel-level probability map $p(x, y) \in [0, 1]$, where each value represents the likelihood of clearcut activity at pixel location (x, y) between the two acquisition dates. This probability map can be interpreted as a *soft clearcut mask* and, after thresholding, converted into a *binary clearcut mask*. These masks can then be post-processed and aggregated into clearcut inferences that provide *Skogsstyrelsen* with actionable insights for downstream analysis, monitoring, and reporting.

RESULTS

An example evaluation on a 64×64 patch is shown in Figure 3. The first two panels show the *before clearcut* and *after clearcut* Sentinel-1 observations (January 2020 and March 2024, respectively). The next panels visualize the model output as a predicted *binary clearcut mask* (obtained by thresholding the predicted probability map at 0.5) alongside the corresponding ground-truth mask. This qualitative example illustrates the end-to-end behavior of the system: given a before/after image pair, the model produces a pixel-level clearcut inference that can be directly compared to reference labels.

Building on this patch-level illustration, we evaluate the clearcut detection model quantitatively across four Swedish regions with varying forest compositions and acquisition geometries: Jönköping County, Västmanland County, Skåne County (with the highest proportion of deciduous trees), and Norrbotten County. For each region, Sentinel-1 image pairs were collected over $50\,\text{km} \times 50\,\text{km}$ tiles; the top half ($25\,\text{km} \times 50\,\text{km}$) of each tile was used for training and the bottom half for validation and testing. Approximately ten image pairs per region were reserved for testing, and twenty random patches per pair were sampled using the Skogsstyrelsen mask.

What the results mean in practice

The evaluation is designed to answer two operational questions:

- *Do we detect most clearcuts of interest?* This is summarized by the True Positive Rate (TPR): how often the system flags a real clearcut (above a minimum size threshold).
- *Do we keep false alarms low?* This is summarized by the False Positive Rate (FPR): how often the system raises a clearcut alarm where no clearcut is recorded (above the same size threshold).

In addition, we report Foreground Mean Intersection over Union (Mean IoU) as an indicator of how well predicted clearcut shapes overlap with reference shapes on a patch basis. While IoU is a technical measure, higher values generally correspond to better spatial agreement and cleaner delineation of clearcut boundaries.

Performance is reported for clearcuts larger than $0.4\,\text{ha}$, which focuses the evaluation on harvest sizes that are most relevant for operational monitoring. It is also important to note that the Skogsstyrelsen annotation masks are based on administrative (optical or manual) records and are not expected to be perfectly aligned with the physical event in all cases. In practice, there can be temporal uncertainty in the recorded cut date and occasional omissions, which can make some real changes appear as false positives or shift detections across evaluation intervals. This is also a key motivation for the timeline analysis, which estimates when clearcut signals begin to appear in the satellite observations. Thus our solution can be effectively used to understand the real time when the clearcut occured.

Overall Performance

Across the four regions (40 image pairs, 777 patches in total), the model achieves:

- Mean Foreground Mean IoU: **0.3**;
- Mean TPR: **0.85**;
- Mean FPR: **0.1**;
- 624 out of 777 patches have IoU > 0, indicating that a substantial majority of patches show at least partial overlap between prediction and reference clearcuts.

Taken together, these results indicate that the system detects the majority of clearcuts while keeping false alarms at a level compatible with downstream review and operational use.

Region-wise Performance

Table I summarizes region-wise performance. The results remain consistently strong across regions that differ in forest composition and acquisition geometry. In particular, TPR stays high across all evaluated regions, showing that the approach generalizes beyond a single local condition, while FPR stays comparatively low, indicating that the model does not rely on fragile, region-specific cues. Variations in Mean IoU are expected given differences in canopy composition, terrain, and seasonal effects; nevertheless, the observed IoU levels confirm that the model is able to produce spatially meaningful clearcut masks rather than only coarse detections.

Timeline Based Analysis

An important objective for forest agencies is to determine *when a clearcut actually occurred*, rather than when it was

TABLE I
REGION-WISE EVALUATION RESULTS.

Region	Mean IoU	Best IoU	TPR	FPR
Jönköping	0.22	0.66	0.67	0.10
Skåne	0.35	0.78	0.89	0.09
Västmanland	0.295	0.85	0.82	0.10
Norrbotten	0.38	0.84	0.80	0.11

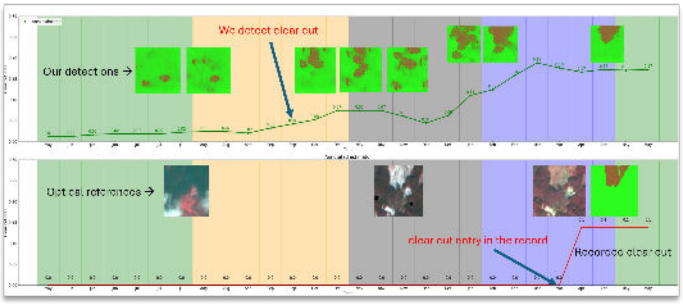

Fig. 4. Example clearcut progression over time, estimated from Sentinel-1 before–after pairs.

later recorded in an administrative database. Our temporal analysis provides a robust, model-driven way to estimate the onset time of a clearcut from satellite observations. For a given region and analysis period, we construct a sequence of *before/after* image pairs that share the same *before* date and vary the *after* date, and then sort these pairs chronologically. For example, if the *before* image is from January, then the pair January–February is the first element in the sequence, January–March is the second, and January–December is the twelfth. Next, we run the trained model on the same region for every image pair in this ordered sequence. For each pair, the model produces a predicted clearcut mask, from which we compute the *clearcut area coverage*. We define this as the fraction of pixels predicted as clearcut:

$$c_t = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}\left(\hat{m}_t(i) = 1\right),$$

where $\hat{m}_t$ is the predicted binary mask for the image pair ending at time t, N is the total number of pixels, and $\mathbb{1}(\cdot)$ is the indicator function. We then plot c_t on the complete timeline. This curve provides an estimate of when clearcut activity begins to appear in the region. When no clearcut has occurred, c_t remains at 0 or a very small value (attributable to model estimation error). Once a clearcut occurs, c_t increases sharply and then typically remains stable until additional harvesting takes place. Because forest regrowth occurs over much longer time scales than our analysis window, the resulting curve is expected to be approximately monotonic over a one-year period: it stays near zero before the event, increases as clearcut area appears, and does not decrease meaningfully within the same year. Figure 4 shows an example of this temporal analysis for the Värmland region in Sweden over the period May 2022 to May 2023. The timeline illustrates how the system captures the progression of harvesting activity over time and can indicate an onset earlier than the date recorded by the forest agency. This earlier detection is consistent with optical reference imagery: the presence of snow-covered exposed ground in the clearing provides visual confirmation of the clearcut at the time indicated by our estimate.

CONCLUSION

This white paper demonstrated an operational clearcut detection approach built on Sentinel-1 SAR and a radar-optimized neural network, enabling monitoring that remains reliable under clouds, darkness, and snow cover. To our knowledge, the work is among the first operational efforts to perform large-scale clearcut detection using Sentinel-1 SAR change pairs and learning-based segmentation as the primary sensing modality. Validated across multiple Swedish regions with differing forest conditions, the system detects the majority of clearcuts while keeping false alarms low for operationally relevant harvest sizes. Despite the known challenges of using C-band Sentinel-1 backscatter for forest change analysis, the achieved results demonstrate that an operational-quality service can be realized when the end-to-end system is optimized around core requirements such as all-weather coverage, cost effectiveness, and long-term reliability. In addition to producing clearcut masks, the timeline analysis illustrates how repeated SAR observations can be used to follow the progression of harvesting activity and estimate when clearcut signals begin to appear, supporting earlier situational awareness than workflows constrained by optical availability. Overall, the results support a scalable, deployment-ready capability for continuous forest monitoring, with straightforward adaptation to new regions through model fine-tuning.