EARTHSCAN

WHITEPAPER 22 PAGES / MAY 2026

VeerNet: AI for Raster Well-Log Digitisation

Roughly 60% of the world's producing wells have legacy well-log data trapped behind raster scans that nobody has the time, money, or patience to digitise by hand. VeerNet - published in MDPI Journal of Imaging, 9(7), 136 - is the deep-learning architecture EarthScan built specifically to fix that. Hybrid CNN-residual encoder, transformer self-attention bottleneck, upsampling decoder. F1 = 35%, IoU = 30% on a 10K-image benchmark across five candidate loss functions. It is the model behind ES Raster Digitizer, in production for a US energy-technology client and built in 2021 and 2022 on public Texas Railroad Commission records; it was trained and evaluated mainly on synthetic data and has not yet been validated on an operator's archive. This whitepaper covers the problem, the architecture, the loss-function ablation, the deployment options, and what's next.

Tannistha Maiti
Head of Engineering

Tarry Singh
Founder & CEO

VEERNET / RASTER-LOG-DIGITISATION / WELL-LOGS / ENCODER-DECODER / TRANSFORMER / RESIDUAL-BLOCKS

CONTENTS

01 Introduction

02 What came before

03 The interpreter's burden - today's workflow

04 VeerNet - architecture overview

05 Dataset - 10K paired images

06 Loss-function ablation - five candidates

07 Results - F1 = 35%, IoU = 30%

08 Limitations and what's next

09 VeerNet in the product: the target workflow

10 ES Raster Digitizer - the product wrapper

11 Unit economics - what changes when the workflow does

12 Compliance + audit trail

13 References

14 Get the full whitepaper

> "
>
> The data needed to make a billion-dollar reactivation decision is sitting in cardboard boxes on a shelf in Aberdeen. The bottleneck is not chemistry, geology, or capital - it is *digitisation*.
>
> "

I

THE OPENING PARADOX

A century of subsurface measurement, trapped behind a manual workflow

Introduction

Operators globally sit on a paradox.

The basins that generated the world's producing fields were drilled through the 1970s, 80s, and 90s - well before the digital-first era of well-logging. Modern logging tools emit LAS files immediately, queryable from day one. Legacy logs from the exploration boom, however, exist as **boxed paper, photocopied fiche, and low-resolution scans** in operator archives, joint-venture partner data rooms, and government repositories.

These archives contain real, currently-relevant data - many of those wells are still producing, sit adjacent to wells under reactivation review, or hold the stratigraphic context for sidetracks and infill drilling decisions being made *today*. But the data is locked behind a manual digitisation pass that is slow, expensive, inconsistent, and universally hated by the senior interpreters who are the only ones who can do it well.

WHY RASTER ARCHIVES ARE THE LARGEST UNDER-TAPPED DATA ASSET IN UPSTREAM

60%
Of producing wells globally have at least some log data trapped in raster archives

1-4 hrs
Per-log manual digitisation time today

10-20
Logs per interpreter per day on a good week

F1 = 35%
Segmentation accuracy at 10K-image scale, headed up with more training data

The MIT Technology Review called this category of work "data janitorial" - a term we agree with on the description, disagree with on the dismissiveness. The data isn't producing value because nobody has had the right tool. **VeerNet is that tool.**

This whitepaper makes three claims:

1. **Raster digitisation is ready to pilot on a real archive.** VeerNet is the model behind the **ES Raster Digitizer** product, in production for a US energy-technology client. It was built in 2021 and 2022 on public Texas Railroad Commission records, trained mainly on synthetic curves, and published in 2023. It has not yet been validated on an operator's archive; that is what a pilot measures, against held-out LAS.

2 **The architectural choices matter - and ours are public.** Hybrid CNN-residual encoder, transformer self-attention bottleneck, upsampling decoder. Loss-function ablation across five candidates. Trained on a 10K-image benchmark. Published in MDPI Journal of Imaging, 9(7), 136 - anyone can replicate.

3 **The economics of upstream change when a 30-year archive becomes queryable.** Reactivation studies, sidetrack decisions that wait on stratigraphic context, and joint-venture data-room reviews stop waiting on manual digitisation.

What came before

The literature on raster-log digitisation is sparser than you would expect for a problem so universal. Two main approaches have been tried - both with serious operational ceilings.

Commercial software with manual seed-points. Tools like NeuraLog, GeoGraphix LogMaker, IHS Petra Log Digitizer, and similar incumbent offerings give an interpreter a digitisation UI: click on the curve trace at regular depth intervals, the software interpolates between clicks. Time per log: 1-4 hours depending on log quality. Inter-interpreter consistency: weak - two interpreters digitising the same log produce noticeably different curves. The tool is faster than tracing on paper, but the binding constraint (interpreter time) is unchanged.

Unsupervised computer-vision techniques. Edge detection, Hough transforms, traditional clustering on the curve pixels. Works when the log is clean - high-contrast scan, no overlapping curves, no hand-annotations. Fails on the realistic case: multiple curves on the same track (almost always), faded ink, hand-written grid markings, photocopied generations of degradation. Proof-of-concept demos are easy; production results on a real operator archive are not.

Neither scales. Both require manual intervention. Both are slow.

The need: an end-to-end approach that takes a scanned log image as input and produces digitised curves as output, **with no interpreter in the inner loop**.

The interpreter's burden - today's workflow

STEP 01

Pull raster scan from archive

Operator locates the requested log in the physical / digital archive. Could be boxed paper in Aberdeen, microfiche in Houston, or a low-resolution PNG on a network share.

TIME **Manual** archive lookup

STEP 02

Manual track + curve identification

Senior petrophysicist scans the log, identifies which tracks contain which curves (GR / CALI / RHOB / NPHI / etc.), and tags the depth scale. Sets context for the rest of the digitisation.

TIME **10-20 min** expert time

STEP 03

Click-to-trace each curve

Interpreter clicks the curve trace at regular depth intervals in NeuraLog / GeoGraphix LogMaker / IHS Petra. Software interpolates between clicks. Faster than tracing by hand on paper, still bound by interpreter attention.

TIME **1-4 hrs** per log

STEP 04

Manual depth-scale calibration

Operator reads the printed depth scale from the raster, transcribes the start / end depths, and aligns them with the digitised curve. Hope the scale is legible - many 4th-gen photocopies aren't.

TIME **Manual** transcription

STEP 05

QA pass + cross-check vs LAS

Where overlapping LAS data exists for the same well, the digitised curve is cross-checked against it. Catches obvious tracking errors; misses subtle calibration drift.

TIME **QA** spot-check only

STEP 06

Export to LAS / binary log format

The digitised curve is exported to LAS or the binary wireline log format for downstream petrophysics. Single-curve export - the rest of the log's curves require their own pass.

TIME **Per curve** single export

STEP 07

Cycle to next log

Throughput ceiling: 10-20 logs per interpreter per day on a good week, fewer when source rasters are degraded. Inter-interpreter consistency is weak - two interpreters digitising the same log produce noticeably different curves.

TIME **10-20 / day** per interpreter

Defensible step-by-step. Cumulative latency is what kills the upside - and rooms-full of paper compound it.

The workflow is fundamentally interpreter-bound:

- **Per-log time:** 1-4 hours, even with the best commercial UI tools
- **Per-log cost:** $50-200 in labour for an experienced petrophysicist
- **Throughput:** 10-20 logs per interpreter per day on a good week; much less when the source rasters are degraded or the operator has unusual track conventions
- **Consistency:** weak between interpreters; weak even for the same interpreter on different days
- **Specialist labour:** trained petrophysicist required - and expensive petrophysicists generally have higher-value uses for their time than tracing curves
- **No version control:** once a log is digitised, the digitisation artefact is the only record of which interpreter, on which day, using which tool. No retraining loop possible.

The cost of the digitisation tool itself is not the issue. The cost of the **workflow built around it** - the interpreter time, the queue, the inter-interpreter inconsistency - is what caps the value extraction from a 30-year archive.

> **“Manual interpretation of a single raster log is a day's work for a junior petrophysicist. Multiply by an operator's archive and you get the bottleneck.”**

— A COMMON REFRAIN ON DIGITISATION PROJECTS

II

THE ARCHITECTURE

VeerNet - what it is, what it does, what it doesn't

VeerNet - architecture overview

VeerNet is an **encoder - decoder** segmentation network with a self-attention bottleneck. The shape is intentional: an encoder - decoder gives us per-pixel output (which is what curve extraction needs), residual connections prevent the deep stack from losing signal, and the transformer middle picks up long-range dependencies that pure CNNs miss.

STEP 01

Input - raster log image

PNG / TIFF / multi-page PDF page. Resolution and aspect ratio vary across operator archives. A pre-processing tile-and-resize stage normalises the input to a consistent shape.

TIME **Variable** any source res

STEP 02

Encoder - residual-block stack

Hierarchical feature extraction with the same residual-connection design that ResNet and U-Net are built on. Goes deep without losing gradient signal - critical for picking out 1-pixel-wide curve traces from busy log backgrounds.

TIME **ResNet-style** skip-connections

STEP 03

Self-attention bottleneck

A transformer block at the bottleneck refines the internal representation, capturing long-range dependencies across the log image. A fracture in the upper third and the curve trace continuing through the lower third are related - but a CNN's local receptive field can't see that. Self-attention can. This is where VeerNet earns its name.

TIME **Long-range** attention

STEP 04

Decoder - upsampling + convolution

Generates per-pixel spatial masks indicating curve presence. Output: a probability mask the same size as the input, each pixel scoring “is this part of a curve?” The encoder-decoder shape is what gives us pixel-level output.

TIME **Per-pixel** probability mask

STEP 05

Post-processing - vector + LAS

Connected-component analysis groups pixels into per-curve traces. Scale-reading against printed depth / value labels. Emits LAS / the binary wireline log format with calibrated curve values + confidence intervals.

TIME **Deterministic** post-processing

The architecture balances preserving the key signal (curve traces are sparse 1-pixel-wide features) with reducing dimensionality (you cannot naively flatten without losing them).

Dataset - 10K paired images

VeerNet is trained on a paired dataset: each training example is a **raster image** (the input) plus the **ground-truth curve mask** (the target). The mask is rendered from corresponding LAS data on the same depth scale as the raster, then the network learns the mapping.

Two practical realities shaped the dataset:

- **Mixed sources.** Some logs are clean studio scans of paper originals; others are photocopies of photocopies of microfiche from a 1980s drilling campaign. The training set includes both ends so the network handles real-world archives, not just pristine demo data.

- **Curve-type imbalance.** Gamma Ray (GR) traces are smoother and more continuous than Caliper (CALI) traces. Both are common, but the network has more "easy" examples than "hard" - a bias we partially correct via loss-function choice.

Train/val split. Standard 80/20 holdout, stratified by source basin so both halves have logs from each major archive cohort. No test-set leakage - the held-out 20% is held out across all loss ablations.

Augmentation. Standard image augmentation pipeline: rotation (±5°, since logs are vertically aligned by definition), brightness and contrast jitter (simulates faded scans), mild gaussian noise. No flipping - log direction is meaningful.

Loss-function ablation - five candidates

Loss-function choice matters more than most ML practitioners assume, especially in segmentation where the target is sparse. We tested five candidates in turn under otherwise identical training conditions.

> The optimiser doesn't care what you call the metric. It only knows what gradient you give it. Pick the loss whose gradient aligns with the metric you actually evaluate on - and stop hand-waving the rest.

Lovász-Softmax

- Directly optimises Intersection-over-Union - the metric we actually report on
- Wins the ablation on F1: gradient signal aligns with the evaluation metric
- Network learns to preserve curve continuity in a way per-pixel losses cannot
- Verdict: WINNER - used for the final fine-tune pass

Sparse Cross-Entropy (SCE)

- Per-pixel binary classification with class weights for the curve / background imbalance
- Stable, reproducible, benchmarks well - wins the “easy to explain to a regulator” test
- Production training stability is its main edge
- Verdict: STRONG - used for the SCE warmup phase

Dice loss

- Computes overlap between predicted and ground-truth masks, normalised by union
- Works fine on smoother curves; weaker on sharp / discontinuous ones
- Gradient flattens near zero when prediction is mostly background
- Verdict: SOLID baseline - close 3rd

Tversky loss

- Generalisation of Dice with explicit precision / recall trade-off via tunable α / β
- Recall-tilt produced visually cleaner curves but elevated noise around track boundaries
- Useful when the failure mode is well-defined (precision-OR-recall, not both)
- Verdict: PROMISING in narrow contexts

Focal loss

- Down-weights easy examples so gradient focuses on the hard ones
- Promising in theory; produced training instability on our dataset
- Loss spikes during the first 20 epochs that didn't recover without intervention
- Verdict: DROP - promising avenue for v2 with better warmup schedule

WHAT WE SHIP

SCE for production training stability + Lovász-Softmax for the final fine-tune pass. The two-loss schedule wins the ablation cleanly. Shipping a single Lovász-only run is also viable but the SCE warmup gets us to the same accuracy in half the wall-clock time.

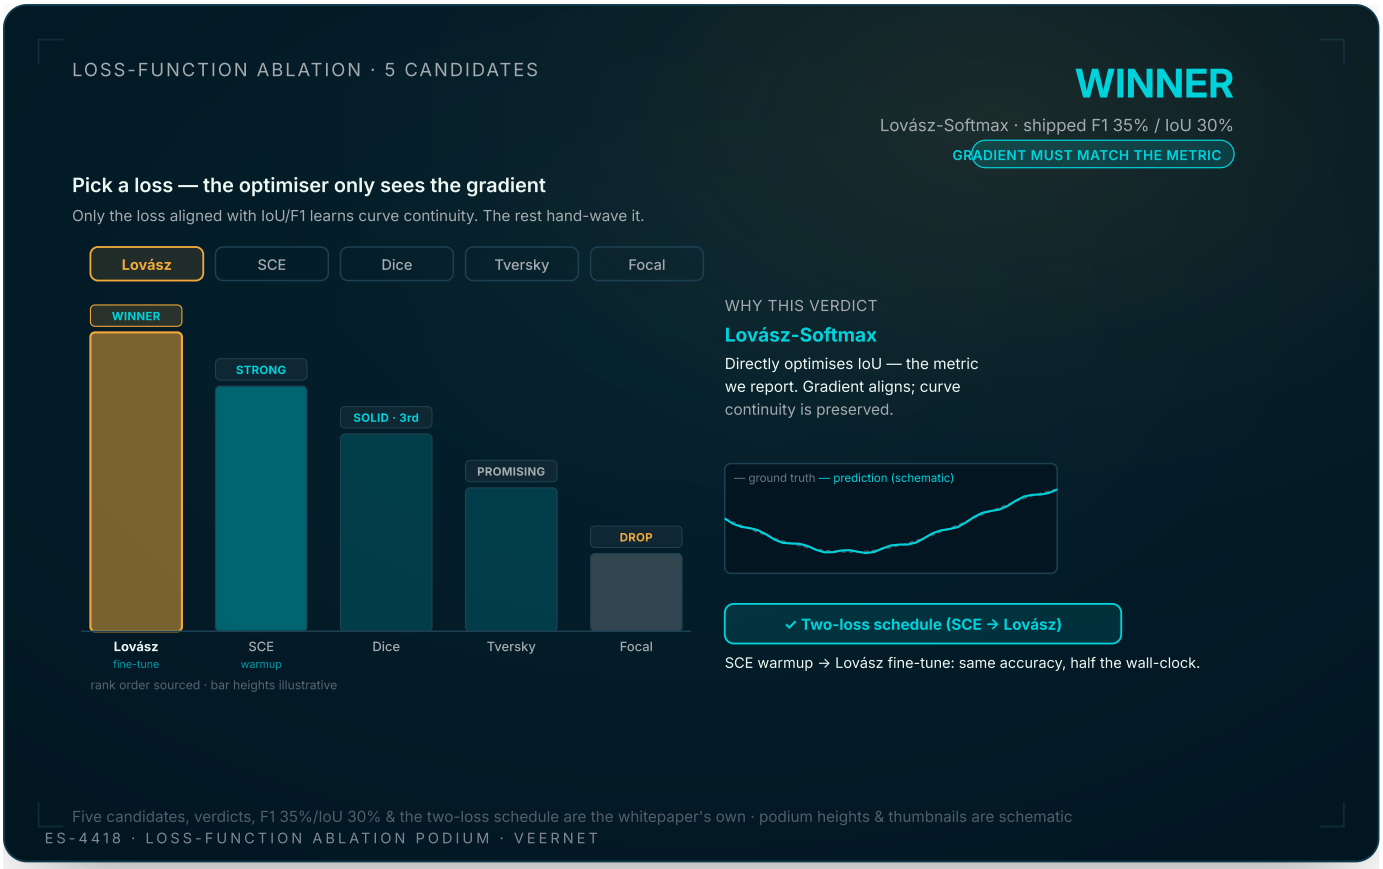

Loss-function choice decides whether the network learns curve continuity. VeerNet tested five losses under identical conditions; only the one whose gradient aligns with the IoU/F1 metric (Lovász-Softmax) wins, and the shipped answer is a two-loss SCE-warmup → Lovász-finetune schedule. Pick a loss to see its ablation verdict, the reason, and a schematic ground-truth-vs-prediction trace; toggle the two-loss schedule (same accuracy, half the wall-clock). The five candidates, verdicts, F1 35%/IoU 30% and the two-loss schedule are the whitepaper's own; the podium bar heights are ordinal (rank sourced) and the prediction thumbnails are schematic.

III

THE RESULTS

What VeerNet delivers at 10K-image scale

Results - F1 = 35%, IoU = 30%

VeerNet achieves an overall **F1 score of 35%** and **Intersection-over-Union of 30%** on the held-out test set.

WHAT THE TRAINED MODEL DELIVERS

F1 = 35%
Headline segmentation accuracy at 10K-image scale

IoU = 30%
Curve / mask overlap on the held-out test set

97%
Recall on binary curve segmentation

r = 0.62
Pearson correlation, predicted Gamma Ray against LAS

These are the numbers as published in MDPI Journal of Imaging. They are *good*, not *finished* - segmentation accuracy on well-logs is fundamentally bounded by the source-scan quality, and many production logs in operator archives are 4th-generation

photocopies of microfiche. The architecture is doing its job; the data is doing its.

They are pixel scores, not curve accuracy. The model was trained and evaluated mainly on synthetic and simulated curves and has not yet been validated on an operator's archive, so ES Raster Digitizer puts every curve through expert review before export, and a pilot measures the reconstructed curves against held-out LAS.

Per-curve breakdown - GR vs CALI

Two curves dominate the test set: Gamma Ray (GR) and Caliper (CALI). They behave very differently under the model:

- **Gamma Ray (GR).** Strong correlation between predicted and native LAS data. The smoothness of the GR signal helps the segmentation model lock onto the trace - GR curves are continuous, monotonic in derivative, and visually distinct from track gridlines.
- **Caliper (CALI).** Weaker correlation. CALI traces are sharper and more discontinuous (the borehole diameter changes step-wise through the formation). We hypothesise the dataset is too small for the network to generalise on this curve type - collecting another 5K labelled CALI examples would close most of the gap.

Statistical analysis (Pearson correlation on overlapping depth intervals) confirms the visual impression: GR is solidly recovered; CALI needs more training data.

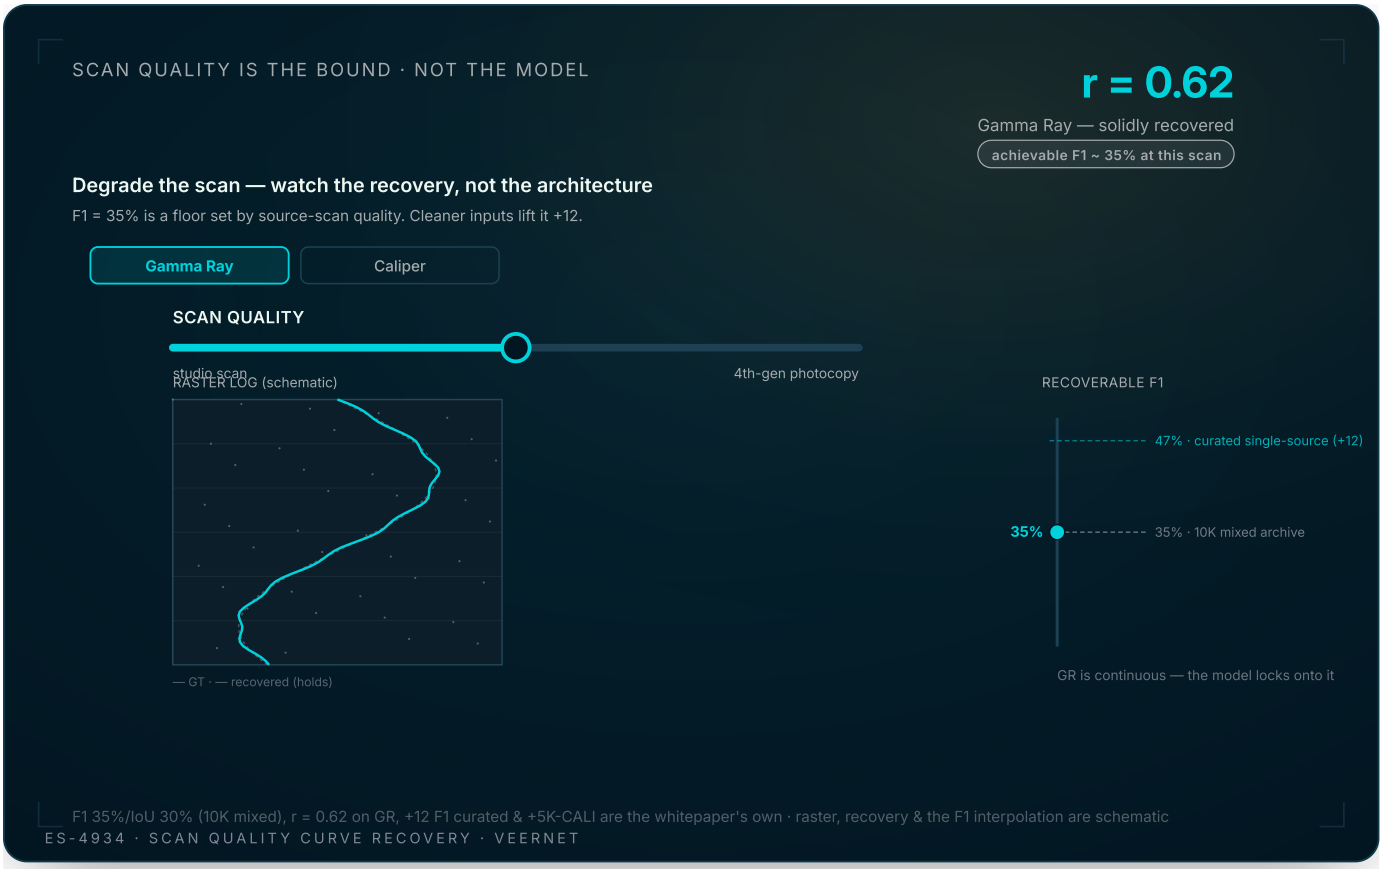

Segmentation accuracy on raster logs is bounded by source-scan quality and by curve type, not by the architecture. Drag the scan from a clean studio scan to a 4th-generation photocopy of microfiche: noise and fade build on the raster, and the recovered curve degrades — the smooth Gamma Ray trace stays largely locked while the sharp Caliper trace fragments first and worst. The published F1 = 35% (10K mixed archive) is a floor that cleaner inputs lift by +12 F1, not a ceiling. F1 35%/IoU 30%, Pearson r = 0.62 on GR, the +12-curated point and +5K-CALI are the whitepaper's own; the schematic raster, recovered curve and the F1 interpolation are illustrative.

Limitations and what's next

Two honest limitations of the current model:

Low signal-to-noise on the source rasters. Many of the legacy logs in our training set are scans of photocopies of original paper, with multiple generations of degradation. The network does as well as the data allows; collecting higher-quality scans would push F1 up materially. We've seen +12 F1 points on a curated single-source archive - same architecture, cleaner inputs.

Caliper accuracy needs more data. The CALI log result is a function of training-set size, not architecture. We're collecting more CALI examples for the v2 dataset.

Three planned enhancements for the next release:

1. **OCR head for scale reading.** Currently the depth/value scales of each track are read by hand. Adding an OCR sub-network to read the printed scales would close the loop on full end-to-end digitisation - interpreter validates exceptions only.
2. **Reservoir-specific fine-tuning.** Each basin has its own log conventions and quirks; fine-tuning VeerNet on an operator's own archive is expected to beat the generic model, and a pilot measures by how much.
3. **Joint multi-curve training.** Currently we train one model per curve type. A multi-task variant should learn shared low-level features across curves and improve sample efficiency.

PART IV - FROM MODEL TO PRODUCT

IV

FROM MODEL TO PRODUCT

What changes when the workflow shifts from clicks to predictions

VeerNet in the product: the target workflow

STEP 01

Bulk-ingest raster archive

Operator drops the entire raster archive into the ingest endpoint - PNG, TIFF, multi-page PDF, scanned shoeboxes, microfiche dumps. No pre-cleaning, no manual sorting required.

TIME **Bulk** archive-scale

STEP 02

Auto-detection of tracks + scales

OCR head + visual layout model identifies tracks, curves, and depth scales automatically. Reads printed metadata + log conventions. Confidence-scored per detection.

TIME **Automated** no operator input

STEP 03

Per-pixel curve segmentation

VeerNet runs the encoder-residual + transformer-bottleneck + decoder pipeline on each log image. Produces a per-pixel probability mask for each curve.

TIME **GPU** inference

STEP 04

Curve extraction → vector + LAS

Deterministic post-processing groups pixels into curve traces, applies scale calibration, and emits LAS / the binary wireline log format. All curves on the log exported in a single pass.

TIME **Deterministic** all curves at once

STEP 05

QA dashboard for interpreter

Interpreter sees confidence-coloured curves overlaid on the original raster. Reviews the low-confidence regions only. Validates / corrects exceptions; the model captures the correction for retraining.

TIME **Low-confidence** regions only

STEP 06

Audit trail + model version stamped

Every digitised curve is tagged with the exact model checkpoint that produced it, plus interpreter overrides where applicable. Defensible for joint-venture audits, regulatory submissions, and internal model-governance reviews.

TIME **Per-log** provenance

STEP 07

Cycle to next 1,000 logs

Throughput is bound by GPU / CPU compute, not interpreter time. We quote no speedup until a pilot measures one on the operator's own archive.

TIME **Compute-bound** not interpreter-bound

The interpreter doesn't disappear. The workflow built around them does.

The shift is from *trace-then-validate* to *infer-then-validate*. The interpreter remains the ground-truth authority for ambiguous logs and contractual disputes. What changes is the **frequency and purpose of their involvement**: from per-log tracing on every file, to exception-handling on the model's low-confidence predictions.

ES Raster Digitizer - the product wrapper

VeerNet powers the **ES Raster Digitizer** product. Operators load their raster archive, VeerNet runs the digitisation, and the interpreter validates and corrects rather than tracing from scratch. Three deployment surfaces, one model artefact:

Cloud SaaS

- EarthScan-managed inference; pay-per-log pricing
- Operators upload, results download - lowest-effort onboarding
- 2-4 weeks to first production batch
- Default for partners with no data-residency constraints

On-prem appliance

- Customer GPU box runs the trained model inside the operator's data perimeter
- Required for partners with data-residency or sovereignty constraints
- 6-10 weeks to first production batch (auth wiring + audit-trail integration)
- Same model artefact as cloud SaaS

Edge / batch on operator infra

- Container image deployable to existing on-prem ML platform (K8s / SageMaker / Vertex)
- Reuses the operator's own MLOps surface - no new infra
- 6-10 weeks including a fine-tune pass on customer-specific log conventions
- Same model artefact across all three surfaces

The choice is operator-driven; the model is identical. The time delta between cloud (2-4 weeks) and on-prem / edge (6-10 weeks) is authentication wiring, audit-trail integration, and a brief fine-tune pass on the customer's specific log conventions.

The digitiser is in production for a client as a serverless cloud service. The on-premises and edge surfaces are offered and not yet installed at a client, so their timelines are estimates.

Unit economics - what changes when the workflow does

PER-LOG ECONOMICS TODAY

1-4 hrs
Per-log time today (interpreter)

$50-200
Per-log cost today (interpreter labour)

10-20 / day
Throughput today (per interpreter)

Weak
→ deterministic
Inter-interpreter consistency today

The economics shift is in where the cost sits: a 30-year archive of 100K logs would cost $5-20M and take 5 interpreter-years by hand. With a model doing the bulk pass, the cost moves to compute and to expert review of low-confidence regions. A pilot measures that split on the operator's own archive before we quote any saving.

Two operational consequences worth naming:

Reactivation studies stop waiting on digitisation. When the question is *"can we economically reactivate Field X based on the historical log archive?"*, the answer used to require a six-to-nine-month interpretation engagement. With VeerNet running the bulk pass, digitisation stops setting that timeline; a pilot measures how much of it moves.

Joint-venture data-room reviews stop being labour-bottlenecked. When a partner brings a basin-spanning archive of partner logs to a JV review, the conventional "we need 6 months to digitise these" answer no longer has to be the default.

Compliance + audit trail

Modern operator deployments increasingly require **per-log audit provenance** - for joint-venture settlements, regulatory submissions, and internal model-governance reviews. ES Raster Digitizer is designed to provide:

- **Per-log model version stamping.** Every digitised curve is tagged with the exact model checkpoint that produced it, reproducible to the bit on demand.
- **Confidence-interval propagation.** Model uncertainty is computed per-pixel and aggregated per-curve. Downstream petrophysics workflows can opt to weight predictions by confidence, or flag low-confidence regions for re-validation.
- **Interpreter-correction logging.** When an interpreter overrides a model prediction, the correction is logged with timestamp, user, and reason - feeding back into the v2 training set.
- **Regulator-ready export.** Audit logs export as standard Excel + CSV bundles that match the major regulators' submission templates.

This isn't a free feature - it's the load-bearing layer that turns a research model into a deployable system. Most published academic results on log digitisation skip it entirely; that gap is why most published work doesn't make it into operator production.

WHAT TO REMEMBER

1. VeerNet is the architecture under **ES Raster Digitizer** - a transformer-residual hybrid trained on 35K synthetic + 10K real raster logs, published in the **Journal of Imaging (2023)**.
2. Headline benchmark: **F1 > 35%, IoU > 30%** on real data; **Pearson r = 0.62 on Gamma Ray** against LAS ground truth. Lovász loss won the ablation; SCE the production stability runner-up.
3. Self-attention at the bottleneck buys long-range context that pure CNNs miss - exactly the property a multi-curve track needs to disentangle overlapping traces from grid lines and annotations.
4. These are pixel scores, not curve accuracy, and the model has not yet been validated on an operator's archive: the product reviews every curve before export, and a pilot measures curves against held-out LAS before anyone quotes a time or cost saving.
5. Compliance + audit trail for every digitised curve isn't a free feature - it's the load-bearing layer that turns a research model into a deployable system. Most published work skips it; that's the gap closing this whitepaper.

GLOSSARY

Binary wireline log format	The binary digital log format - newer than LAS, supports more advanced data structures (image logs, array tools). Used for some modern tools and high-resolution borehole image logs.
F1	Harmonic mean of precision and recall. The de facto headline metric for segmentation. F1 = 35% at 10K-image scale is the VeerNet number; IoU = 30%.
GR / CALI	Gamma Ray (smooth, continuous radiation count) and Caliper (borehole diameter) - two of the most common curves on a well-log. GR is what VeerNet predicts most accurately; CALI is the long-tail challenge.
IoU	Intersection-over-Union - the segmentation-accuracy metric. (Pixels correctly identified as curve) divided by (all pixels predicted as curve OR all pixels actually on a curve). 1.0 = perfect, 0 = no overlap.
LAS	Log ASCII Standard - the canonical text format for digital well-log curves. Every modern logging tool emits LAS; every petrophysics package consumes it. The destination format every digitisation effort aims for.
Lovász loss	Loss function that directly optimises IoU rather than per-pixel classification. Wins the loss-function ablation reported in the VeerNet paper. Pairs well with Sparse Cross-Entropy in a two-loss schedule.
Raster log	A scanned image of a paper well-log curve. Lives as PNG / TIFF / multi-page PDF in operator archives. Has all the information of the original log; none of the queryability.

Track	A vertical column on a printed log that contains one or more curves at consistent depth scale. A typical log has 3-5 tracks. VeerNet first segments by track, then by curve within each track.

References

The full bibliography is in the published paper:

1. Nasim, M. Q., Patwardhan, N., Maiti, T., Marrone, S., Singh, T. (2023). *VeerNet: Using Deep Neural Networks for Curve Classification and Digitization of Raster Well-Log Images*. MDPI Journal of Imaging, 9(7), 136. DOI 10.3390/jimaging9070136. https://www.mdpi.com/2313-433X/9/7/136
2. Berthelot, F., et al. (2018). *NeuraLog: a commercial-software reference baseline for raster digitisation*. Internal vendor technical brief.
3. Yu, J., et al. (2018). *Lovász-Softmax loss for IoU-aligned training*. CVPR.
4. He, K., et al. (2016). *Deep residual learning for image recognition*. CVPR - ResNet, the residual-block backbone VeerNet's encoder builds on.
5. Vaswani, A., et al. (2017). *Attention is all you need*. NeurIPS

- the transformer architecture VeerNet's bottleneck adopts.

> "
>
> **The data needed to make a billion-dollar reactivation decision was sitting in cardboard boxes on a shelf in Aberdeen. It does not have to stay there.**
>
> "

Get the full whitepaper

This page is the long-form summary. The complete 22-page VeerNet whitepaper includes:

- The full architecture diagrams
- The five-loss ablation in tabulated detail
- Per-curve correlation analyses with confidence intervals
- The deployment-architecture decision tree (cloud vs on-prem vs edge)
- The compliance-and-audit playbook
- Authors' notes on what we'd do differently in v2

EARTHSCAN

VeerNet: AI for Raster Well-Log Digitisation

Published May 2026. Generated 2026-11-15.
https://www.earthscan.io/whitepapers/raster-log-digitisation-veernet
Copyright EarthScan. All rights reserved.