

Quality Log

Secondary porosity in a carbonate reservoir is usually reported as a single vug percentage per interval, a scalar that two very different rocks can share. We describe a different deliverable, built and validated with a mid-sized Middle East carbonate operator: a continuous per-interval morphology log. Every 2 m of borehole yields a six-track statistical fingerprint - a kernel-density estimate of vug area (observed 1 to 12 square centimetres, mass concentrated between 1 and 3.5 and predominantly 1 to 6), a kernel-density estimate of circularity (0.28 to 0.85, peaking semi-circular between 0.45 and 0.7), a four-quadrant azimuth spectrum (0 to 90, 90 to 180, 180 to 270, 270 to 360 degrees), a vug count, the vug percentage itself defined as vug area over interval area, and an estimated-versus-manual overlay track that validates the automated estimate interval by interval. Stacked down the well, these tracks are a log in the petrophysical sense: you read them vertically, you spot the interval where the character changes, and you correlate that signature to the same signature in the next well. The percentage is one of the six tracks, not a summary of the other five. We show that two intervals carrying an identical percentage separate cleanly in the area and circularity tracks and in the azimuth spectrum, that reservoir-quality populations occupy distinct regions of the area-circularity plane a scalar cannot see, and that well-to-well correlation - a Phase-3 objective of the programme - is tractable on the multi-track signature and ambiguous on the percentage alone, because many intervals repeat the same percentage. The statistical outputs are computed on a fine grid (per 10 cm: total count, total area, mean area, standard deviation of area, and the area, circularity and azimuth spectra), and the whole log is produced without a human picking individual pores. The argument is not that percentages are wrong. It is that a percentage is a projection, the fingerprint is the data, and a reservoir-quality workflow should consume the data.

Tannistha Maiti

Senior AI Researcher

Quamer Nasim

ML Research Engineer

VUG-MORPHOLOGY / SECONDARY-POROSITY / CARBONATE-RESERVOIR / RESERVOIR-QUALITY / BOREHOLE-IMAGE-LOGS / MICRORESISTIVITY-IMAGING

CONTENTS

01	The scalar and the curve it came from
02	Six tracks, one depth axis
03	The overlay track earns the trust
04	The shape plane a percentage cannot see
05	Correlation is where the log pays for itself
06	What the statistics underneath actually are
07	Turning the tracks into a pore-type facies log
08	Why the log costs no more than the number
09	Where the fingerprint belongs in the workflow
10	Limitations
11	References

A petrophysicist does not read a gamma-ray value. She reads a gamma-ray log: a curve down depth, whose shape is the information. A single value at a single depth tells her almost nothing; the way the curve climbs into a shale, holds flat through a clean sand, and spikes at a hot streak is what she interprets. The vertical continuity is the point. That is what a log is.

Vug porosity, oddly, is usually not reported that way. The convention in carbonate interpretation is a number: this interval is 6 percent vug. It is a value, stripped of the curve it came from. The reservoir-quality specialist who receives it has to trust that the number stands in for the rock, and in a carbonate it very often does not, because two intervals at the same percentage can be built from completely different pore populations that store and flow fluid completely differently.

This whitepaper is about a deliverable we built with a mid-sized Middle East carbonate operator that puts vugs back into log form. Instead of one number per interval, the pipeline emits a six-track statistical fingerprint every 2 m: the distribution of vug area, the distribution of circularity, the azimuth spectrum, the count, the percentage, and an overlay that plots the automated estimate against the manual interpretation. Stacked down the borehole, those six tracks are a log. You read them vertically, you find the depth where the pore character changes, and - this is the part a percentage cannot do - you correlate that signature to the next well. It is written for the petrophysicist and reservoir engineer who consume vug output, and for the R&D lead deciding what shape that output should take.

The scalar and the curve it came from

We have argued the narrow version of this point before, and there is no value in re-deriving it: a bulk vug percentage is a lossy projection of a per-vug catalogue, and two intervals at the same percentage flow differently. That case, with the individual-vug quantification pipeline behind it, is made in full in our companion whitepaper on AI-assisted individual vug quantification, and a reader who wants the stage-by-stage image-processing chain (mode subtraction, adaptive thresholding, contour extraction, the area-and-circularity gate, the false-positive filters) should start there.

The claim here is one level up, and it is about form, not just content. Suppose you accept that the per-vug catalogue is richer than the percentage. The question this paper answers is: what do you hand the reservoir team? A table of thousands of individual vugs is not

consumable; a single percentage per interval is consumable but lossy. The fingerprint log is the deliverable in between, and it is the right one, because it is consumable in exactly the way a petrophysicist already knows how to consume a log while preserving the distributional information the percentage throws away.

Concretely, the percentage is defined per interval as vug area over interval area. That definition is not in dispute and we compute it faithfully. Our point is that it is one track of six, computed alongside the others, not a summary that replaces them.

VUG PERCENTAGE: ONE TRACK, DEFINED PER INTERVAL

Formula

LaTeX


Copy LaTeX

$$\text{vug}\% (z) = \frac{\sum_k a_k(z)}{A_{\text{interval}}(z)} \times 100$$

Here the numerator sums the areas of the individual vugs detected in the interval at depth z , and the denominator is the imaged area of the interval. Everything that makes the numerator - how many vugs, how large, how round, which way they trend - is discarded the moment you keep only the ratio. The fingerprint keeps it.

Six tracks, one depth axis

The demonstration well below makes the log concrete. Six tracks share a single depth axis, one fingerprint per 2 m. Read left to right, they are the size distribution, the shape distribution, the orientation spectrum, the count, the percentage, and the validation overlay. Read top to bottom, each is a curve down depth that you interpret the way you interpret any log.



Vug detection read as a continuous well log. Six depth-registered tracks share one axis down a demonstration reservoir interval, one fingerprint per 2 m: area KDE (vug size 1 to 12 cm²), circularity KDE (0.28 to 0.85, peaking semi-circular between 0.45 and 0.7), a four-quadrant azimuth spectrum (0-90, 90-180, 180-270, 270-360 degrees), vug count, vug percentage, and an estimated-versus-manual overlay. Two intervals are marked because they carry the same 6.2% vug porosity yet diverge sharply: pin either one to read its full fingerprint in the side panel and see that Interval A is a few large, oriented vugs while Interval B is many small, isotropic ones. Vug percentage is the one track drawn in orange because it is the scalar the exhibit argues against: a single percentage cannot separate those two rocks, and only the multi-track log can. Because every track is depth-continuous, the fingerprint reads down the well and lines up across wells the way any petrophysical log does. The ranges, quadrants, and the vug-percentage definition are sourced from the final estimation dashboards; the per-interval track values down the demonstration well and the twin pair are illustrative shapes drawn to make the argument legible.

The two marked intervals are the argument. They carry the same 6.2 percent vug porosity, so on the percentage track alone they are indistinguishable - the same bar, at two depths. Pin either one and the fingerprint pulls them apart. One is a handful of large, rounded, connected vugs with a strongly oriented azimuth; the other is a swarm of small, angular, isolated pores with a near-uniform azimuth spectrum. Those are not two readings of the same rock. They are different rocks with different storage geometry and, when the large vugs connect, different flow, and the only reason the specialist can tell them apart is that the log preserved the tracks the scalar collapsed.

It is worth being precise about what each track is, because the precision is what makes the tracks trustworthy rather than decorative.

The **area track** is a kernel-density estimate of vug size for the interval, not a single mean. On this reservoir the size distribution runs from 1 to 12 square centimetres, with the mass concentrated between 1 and 3.5 square centimetres in one well and between 1.5 and 4 in another, and predominantly between 1 and 6 across the section. A right-skewed distribution with a tail to large connected vugs reads very differently from a tight distribution of small pores, even when both integrate to the same total area, and the KDE shows that difference at a glance where a mean would hide it.

The **circularity track** is a kernel-density estimate of shape, where circularity is the ratio of a vug's contour area to the area of its minimum enclosing circle, normalised to run from 0 to 1. On this reservoir circularity spans 0.28 to 0.85 and peaks in a semi-circular band between 0.45 and 0.7. Shape is a proxy for pore type and connectivity: rounded, high-circularity vugs behave differently from angular, low-circularity ones, and the track carries that where the percentage cannot.

The **azimuth track** is a four-quadrant spectrum - the mass of vug orientation in the 0 to 90, 90 to 180, 180 to 270, and 270 to 360 degree bins. An interval whose vugs cluster in one quadrant has a directional fabric; an interval whose mass spreads evenly across all four is isotropic. That distinction matters for permeability anisotropy and for anyone building a fracture-vug story, and it is invisible to a bulk number.

The **count track** and the **percentage track** are the two scalars, kept because they are genuinely useful and genuinely consumed, not discarded out of purism. The point is only that they sit beside the distributions rather than in place of them.

The overlay track earns the trust

A morphology log that no one trusts is a curiosity. The sixth track exists to answer the reviewer's first question: does the automated estimate agree with the interpreter? It plots the estimated vug percentage against the manual interpretation-software percentage, interval by interval, as two traces on the same depth axis. Where they track closely, the automated log stands on its own; where they diverge, the divergence is localised to an interval a human can pull and check, rather than hidden inside a single well-level number.

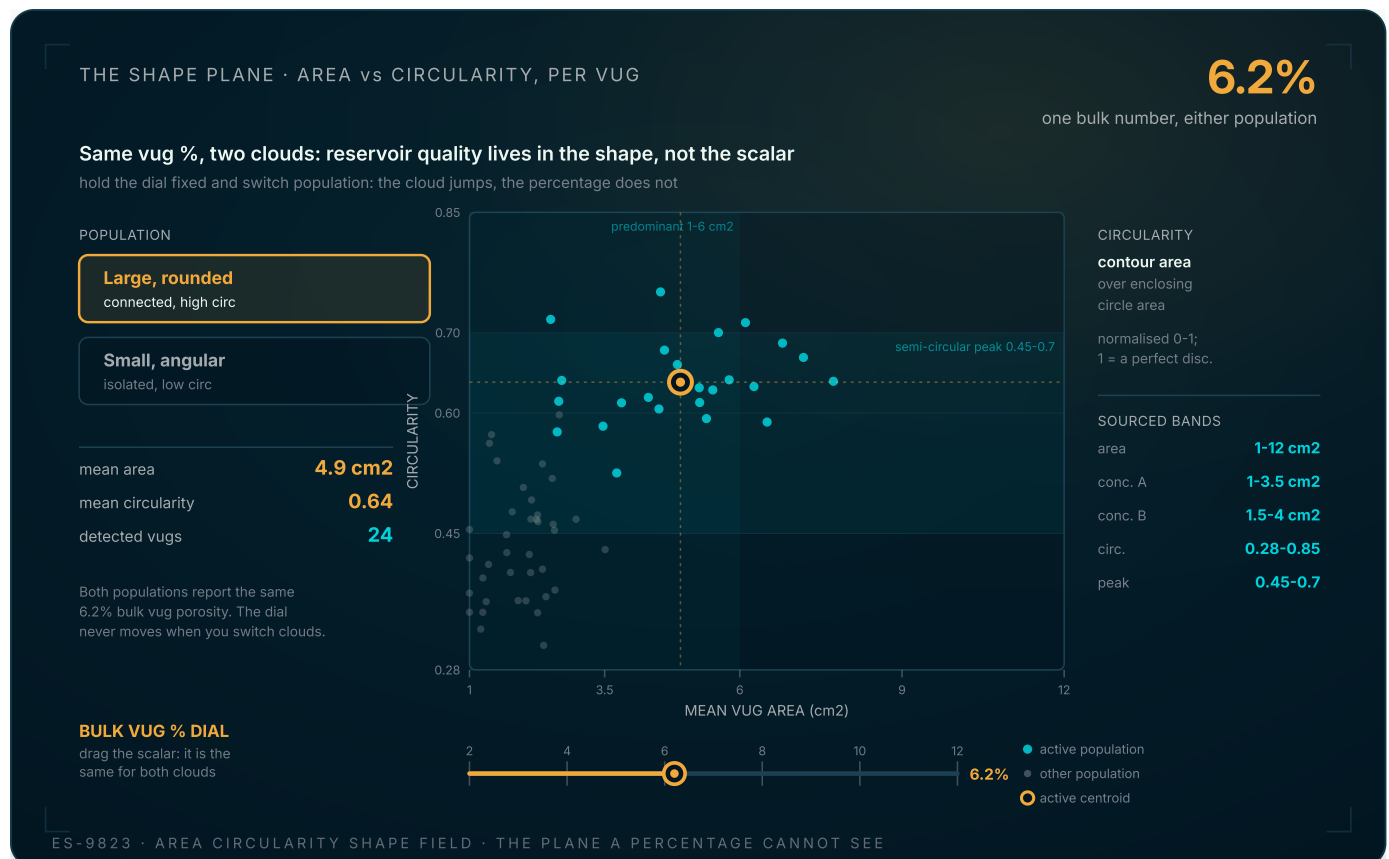
This is the discipline that separates a research figure from a deliverable. The estimated-versus-manual overlay is not a validation you run once and file. It is a permanent track of the log, so every interval carries its own agreement evidence, and a reservoir engineer reading the log three years later can see, at the depth she cares about, whether the automated pick and the human pick agreed. When they do, over dozens of consecutive intervals, the case for reading the automated fingerprint instead of re-picking by hand is made on the log itself, not in a slide.

WHY THE OVERLAY IS A TRACK, NOT A FOOTNOTE

A single well-level agreement statistic tells you the log is roughly right on average. A per-interval overlay tells you exactly where it is right and exactly where to look when it is not - which is the only form of validation a reservoir team can act on interval by interval.

The shape plane a percentage cannot see

The six-track log reads vertically, but there is a second way to read the same fingerprint that makes the "same percentage, different rock" fact undeniable, and it is worth showing directly. Take every individual vug in an interval and plot it in the plane of mean area against circularity. That plane is where reservoir-quality populations separate.




The shape plane a bulk percentage cannot see. Every detected vug of a demonstration interval is plotted by mean area (cm², x) against circularity (contour area over minimum-enclosing-circle area, 0 to 1, y). The sourced reference bands are shaded: the predominant size band 1 to 6 cm² and the semi-circular circularity peak 0.45 to 0.7. Two reservoir-quality populations occupy visibly separate regions: large, rounded, connected vugs sit high-area and high-circularity, while small, angular, isolated pores sit low-area and low-circularity. The dial at the bottom sets the bulk vug percentage, and the exhibit's whole point is that it stays fixed while you switch populations: the cloud jumps across the plane, the number does not. The orange crosshair marks the active population's centroid, the one element that argues. The area and circularity ranges and the area-ratio circularity definition are sourced from the final estimation dashboards; the individual scatter points and the bulk-percentage dial are illustrative.

Hold the bulk percentage fixed on the dial and switch between the two populations. The cloud jumps across the plane - large rounded vugs to the high-area, high-circularity corner, small angular pores to the low-area, low-circularity corner - and the number on the dial never moves. That is the whole problem with the scalar in one gesture. Two rocks that a petrophysicist would never confuse, that a core analyst would classify into different pore types, that flow differently in the reservoir, report the same percentage. The area-circularity plane resolves them; the percentage cannot, because a percentage is a one-dimensional shadow of a two-dimensional (and, with azimuth, three-dimensional) distribution.

The sourced bands anchor the plane to the real reservoir rather than to a generic illustration. The predominant size band of 1 to 6 square centimetres and the semi-circular circularity peak of 0.45 to 0.7 are where most of this reservoir's vugs actually live, and the populations we contrast sit inside those bands. The separation is not an artefact of exaggerated inputs; it is the geometry of two ordinary carbonate pore populations that a bulk number happens to merge.

Correlation is where the log pays for itself

The strongest argument for the fingerprint is not that it describes one well better. It is that it correlates two wells, and a percentage cannot. Well-to-well correlation of vug character was an explicit Phase-3 objective of the programme, and it is exactly the task where the difference between a scalar and a signature stops being philosophical and starts being operational.



Correlating two wells on morphology instead of on a percentage. Each well is a stack of per-2m fingerprints, drawn as a compact three-channel signature: a teal bar for the area-KDE peak (1 to 12 cm²), a bright tick for the circularity peak (0.28 to 0.85), and a chip for the dominant azimuth quadrant. Pick an interval in Well 1 and the exhibit ranks every interval in Well 2 by signature distance, marks the single best match, and draws the orange tie-line. Toggle the correlation basis from the full fingerprint to vug-percentage-only, and the best match jumps to an ambiguous interval because many intervals share the same percentage: the count in the corner is how many right-well intervals tie for best under the chosen metric, and it collapses to one only when the full fingerprint is used. The channels, ranges, quadrants, and the well-to-well correlation task are sourced from the programme; the specific per-interval signature values in each demonstration well and their depth registration are illustrative shapes drawn to make the argument legible.

Pick an interval in the first well and the tool ranks every interval in the second well by how close its fingerprint is - area, circularity, and dominant azimuth together - and marks the single best match. Toggle the correlation basis to vug-percentage-only and the match falls apart: because many intervals in a well repeat the same percentage, the "best match" becomes a tie among several depths, and the tie count in the corner jumps from one to many. A signature built from three morphology channels is specific enough to line up one interval against one interval. A single percentage is not; it is a value that recurs, and a recurring value cannot anchor a correlation.

This is the practical difference. A reservoir team correlating flow units across a field is asking, at each depth in a new well, "which interval in the wells I already understand does this look like?" The fingerprint answers that question because the answer is a shape, and shapes are distinctive. The percentage answers "which intervals share this number?" and the answer is usually "too many to be useful." The log correlates because it is a log; the number does not because it is a number.

What the statistics underneath actually are

The fingerprint is not hand-waving over a picture. Beneath every 2 m track sits a fine-grained statistical computation, run on a 10 cm grid down the whole well, and it is the same set of outputs for every interval, which is what makes the tracks comparable across depth and across wells. Per 10 cm the pipeline computes the total vug count, the total vug area, the mean vug area, the standard deviation of area, and the full area (or porosity), circularity, and azimuth spectra. The 2 m tracks in the log are aggregations of that finer grid, so an interval that looks anomalous can be drilled into at 10 cm resolution without recomputing anything.

That regularity is what turns a pile of per-vug measurements into a log. A log is only useful if every depth reports the same channels in the same units, so that the curve means the same thing at 2650 m as at 2700 m and in well one as in well two. The per-10 cm statistical contract is the mechanism that guarantees it. Two facts follow from it that matter to anyone deploying this. First, the log is reproducible: the same well produces the same fingerprint because the statistics are deterministic functions of the detected vugs, not judgement calls. Second, the log is auditable: every track value traces back to a set of individual vugs with areas and circularities, so a disputed interval can be opened all the way down to the contours that produced it.

Turning the tracks into a pore-type facies log

Once the fingerprint is regular down depth, a second deliverable falls out almost for free: a vug facies log. The reservoir geologist rarely wants the raw distributions at every metre; she wants to know where the interval changes character - where the log switches from a large-connected-vug facies to a small-isolated-pore facies, because that switch is a flow-unit boundary. The multi-track fingerprint supplies exactly the axes that separate those

facies. Large area with high circularity and a directional azimuth is one facies; small area with low circularity and an isotropic azimuth is another; a moderate, mixed signature is a third. Because each track is a curve, the boundaries between facies are places where several tracks turn at once, and a boundary that several independent channels agree on is a boundary a geologist can defend.

This is the difference between the fingerprint and a percentage made vivid a third way. A percentage log can only draw a boundary where the number steps up or down, and the number steps for reasons that mix size, count, and shape into one figure - so a facies boundary and a mere change in vug count look identical on it. The fingerprint separates those causes. An interval where the count rises but the area distribution and circularity hold steady is more of the same rock; an interval where the area distribution shifts to larger, rounder vugs while the count falls is a genuine change in pore type. The percentage cannot tell those two apart because it collapses count and size into a single ratio. The tracks keep them on separate axes, and separate axes are what let you cut facies honestly.

The facies log is also where the azimuth track stops being a curiosity and starts being load-bearing. A directional pore fabric - vug mass concentrated in one or two quadrants - is a permeability-anisotropy signal, and anisotropy is a property a static or dynamic reservoir model must be told about explicitly. On a percentage log that signal does not exist; there is nowhere for it to live. On the fingerprint it is a track, so the interval where the fabric turns from isotropic to strongly oriented is visible, correlatable, and available to hand to whoever is building the permeability model. The morphology log is, in that sense, not just a better description of vugs. It is a set of inputs a downstream model can consume that a percentage never could produce.

Why the log costs no more than the number

The obvious objection to any richer deliverable is cost: a six-track log sounds like six times the work of one percentage. It is not, and the reason is worth stating plainly because it changes the build-versus-consume calculation. The expensive part of vug quantification is the detection - finding and outlining every individual vug in the image, with a false-positive filter disciplined enough to reject the conductive bedding planes and fractures a naive thresholder mistakes for pores. That work is done once, and it is the same work whether the output is a percentage or a fingerprint. The percentage is a single reduction of the

detected vugs; the fingerprint is a handful of reductions of the same detected vugs. The detected vugs are the cost. The reductions are cheap arithmetic over a set you already have.

So the choice between shipping a number and shipping a log is not a choice about how much to compute. It is a choice about how much of what you already computed to throw away before you hand it over. The percentage throws away the area distribution, the circularity distribution, the azimuth spectrum, and the count, and keeps a single ratio. The fingerprint keeps them. Given that the detection has already paid for all of it, keeping it is the default that needs no justification, and discarding it is the choice that does.

There is a reproducibility dividend on top of the cost argument. Because every track is a deterministic statistic over the detected vugs, and because the detection runs the same way on the same image without a human picking individual pores, the log is the same every time the well is processed. Two analysts do not produce two different vug logs the way two interpreters produce two different manual percentages; the inter-interpreter variance that plagues manual vug picking is designed out. A log that is identical on re-run, auditable to the contour, and validated interval by interval against the manual pass is a different class of deliverable from a hand-picked number, and it costs the operator nothing beyond the detection they were going to pay for regardless.

Where the fingerprint belongs in the workflow

The morphology log is not a replacement for the petrophysicist, and it is not a replacement for the bulk percentage where a bulk percentage is genuinely all that is needed - a quick-look net-pay screen, say. It is the deliverable for the questions the percentage cannot answer, and in a carbonate those questions are the ones that decide reservoir quality: what pore type dominates this interval, how does the pore population change down the well, does this flow unit continue into the next well, is there a directional fabric that a permeability model has to honour.

For an R&D lead or a chief geoscientist evaluating whether to consume vug output as a log rather than a table, the decision criteria are concrete:

- Does the deliverable preserve the distribution of vug area and circularity, or collapse it to a mean or a percentage?

- Does it carry orientation as a spectrum, so that a directional pore fabric is visible rather than averaged away?
- Is every automated interval validated against the manual interpretation on the log itself, so agreement is legible at the depth of interest?
- Is the underlying statistical grid fine and regular enough (here, per 10 cm) that any interval can be audited to the individual vugs that produced it?
- Can two wells be correlated on the signature, or only compared on a scalar that recurs and therefore cannot anchor a match?

A log that answers yes to these is a different object from a percentage. It costs no more to produce - the individual vugs are already detected, and the tracks are aggregations of statistics the pipeline already computes - and it hands the reservoir team the curve, not the shadow of the curve. The percentage is still there, as one of the six tracks. It is simply no longer pretending to be the whole log.

WHAT THIS WHITEPAPER ARGUES

- 1 A vug-porosity percentage is a lossy scalar; the deliverable that carries reservoir quality is a continuous per-interval morphology log, one six-track statistical fingerprint every 2 m: area KDE (1 to 12 cm², predominant 1 to 6), circularity KDE (0.28 to 0.85, semi-circular peak 0.45 to 0.7), a four-quadrant azimuth spectrum, count, percentage, and an estimated-versus-manual overlay.
- 2 The percentage (vug area over interval area) is one track of six, computed alongside the distributions, not a summary that replaces them - two intervals at an identical 6.2 percent separate cleanly in the area, circularity and azimuth tracks.
- 3 The estimated-versus-manual overlay is a permanent track, not a one-off validation, so every interval carries its own agreement evidence and a reviewer can see exactly where the automated log and the human pick agree and where to look when they do not.
- 4 Reservoir-quality populations occupy distinct regions of the area-circularity plane that a bulk percentage cannot see: hold the percentage fixed and the pore cloud still jumps from large-rounded to small-angular, because the scalar is a one-dimensional shadow of a multi-dimensional distribution.
- 5 Well-to-well correlation, a Phase-3 objective, is tractable on the multi-track fingerprint and ambiguous on the percentage: a three-channel signature lines up one interval against one interval, while a single percentage recurs across many intervals and cannot anchor a match.
- 6 The tracks rest on a regular per-10 cm statistical contract (total count, total area, mean area, standard deviation of area, and the area, circularity and azimuth spectra), which makes the log reproducible, comparable across depth and wells, and auditable down to the individual vug contours.

Limitations

The morphology log inherits the limits of the detection that feeds it, and honesty about them is part of the deliverable. Borehole image logs capture a two-dimensional cross-section of a three-dimensional pore, so area and circularity are measured on the trace the borehole wall cuts through a vug, not on the vug itself; a large vug clipped near its edge

and a small vug imaged through its middle can present similar traces, and the fingerprint cannot always distinguish them without proximity or graph-based analysis we have flagged as future work. The image resolution sets a hard floor on depth registration - one pixel of the digital log format corresponds to about 3 cm of depth, so the 2 m and 10 cm grids are quantised against that floor and very fine structure below it is not resolved. The estimated-versus-manual overlay validates against an interpretation-software percentage that is itself a human, subjective product with a known bias toward missing pores, so agreement with it is evidence of consistency with expert practice, not agreement with a physical ground truth, and the intervals where the automated log recovers vugs the manual pass missed are precisely the intervals where "disagreement" is a feature rather than an error. The specific per-interval track values, the twin intervals, and the two demonstration wells shown in the exhibits are illustrative, seeded shapes drawn to make the argument legible; the ranges, bands, quadrants, the per-10 cm statistical contract, and the vug-percentage definition they fill are sourced from the programme's final estimation dashboards. Finally, the log characterises vug morphology; it does not by itself resolve connectivity or effective permeability, which need a pore-network or flow model the fingerprint is an input to, not a substitute for.

References

Lucia, F. J. Carbonate Reservoir Characterization: An Integrated Approach. Springer, second edition (2007). Reference on vuggy secondary porosity, pore-type classification, and the control of vug morphology and connectivity on carbonate reservoir quality.

<https://doi.org/10.1007/978-3-540-72742-2>

Suzuki, S., and Abe, K. Topological Structural Analysis of Digitized Binary Images by Border Following. Computer Vision, Graphics, and Image Processing (1985). The contour-tracing algorithm that produces the individual vug outlines the per-interval statistics are computed over. [https://doi.org/10.1016/0734-189X\(85\)90016-7](https://doi.org/10.1016/0734-189X(85)90016-7)

Rosenblatt, M. Remarks on Some Nonparametric Estimates of a Density Function. Annals of Mathematical Statistics (1956). The kernel-density-estimation foundation for the per-interval area and circularity distribution tracks. <https://doi.org/10.1214/aoms/1177728190>

Vug Fingerprints, Not Vug Percentages: Per-Interval Morphology as a Reservoir-Quality Log

Published April 2025. Generated 2026-07-22.

<https://www.earthscan.io/whitepapers/vug-fingerprints-not-vug-percentages-per-interval-morphology-log>

Copyright EarthScan. All rights reserved.