

Reservoirs

Secondary porosity in carbonate reservoirs is dominated by vugs, and the way it is reported in practice — a bulk vug percentage per interval, hand-picked on a borehole image log — discards almost all of the information that controls reservoir quality. A 6 percent interval made of a few large connected vugs and a 6 percent interval made of many small isolated ones flow very differently, yet the bulk percentage cannot tell them apart. This whitepaper presents a classical computer-vision pipeline — deliberately not deep learning — that quantifies vugs individually rather than in bulk. The pipeline chains top-K mode subtraction for local-variation enhancement, Gaussian-modulated adaptive thresholding, Suzuki-Abe contour extraction, an area-and-circularity geometric gate, centroid-and-IoU aggregation to de-duplicate, and a two-stage false-positive filter (Laplacian-variance contrast and mean-intensity differencing) that suppresses the conductive bedding planes and fractures a naive thresholder would mistake for vugs. The output is not a percentage but a distribution: per-vug area (observed 1 to 12 square centimetres, concentrated between 1 and 3.5), circularity (0.28 to 0.85, peaking semi-circular at 0.45 to 0.7), and an azimuth spectrum, computed every 0.1 metre. Validated on three wells spanning both vertical and horizontal high-resolution borehole image logs acquired with two different microresistivity imaging tools, using a single shared parameter set, the pipeline matches the expert's interpretation-software ground truth and, at several intervals, recovers dissolution vugs the manual interpreter missed. Because every step is an interpretable operator rather than a black box, the result is a defensible, auditable, per-vug catalogue — and a natural source of ground-truth masks for a future supervised vug-segmentation model. We frame the work as applied computer-vision engineering: the geoscience defines the gate, the algorithm enforces it reproducibly, and the production layer makes it run on two-hundred-metre logs without a human in the loop.

Quamer Nasim

ML Research Engineer

Tannistha Maiti

Senior AI Researcher

Tarry Singh

Founder & CEO

VUG-QUANTIFICATION / SECONDARY-POROSITY / CARBONATE-RESERVOIR /
COMPUTER-VISION / ADAPTIVE-THRESHOLDING / IMAGE-PROCESSING

CONTENTS

01	Why "not deep learning" was the right engineering call
02	The pipeline, stage by stage
03	The false-positive problem, and the two-stage filter that solves it
04	One parameter set, three wells, two tools
05	What the catalogue reveals that a percentage cannot
06	Catching what the interpreter missed
07	The engineering around the algorithm
08	What this is the foundation for
09	References

A vug is a cavity. Dissolution carves it into a carbonate matrix, and on a borehole image log it shows up as a dark, roughly rounded blob where the rock conducts more than its surroundings. Vugs are where a large share of a carbonate reservoir's secondary porosity lives, which means they are where a large share of its storage and — when they connect — its flow lives too. So the obvious question for a reservoir-quality specialist is: how much vug is in this interval, and what shape is it in?

The standard answer collapses that question into a single number. An interpreter loads the image log into the interpretation software, marks the vuggy zones by hand, and reports a vug percentage per interval — vug area over interval area. It is a reasonable summary and it is also a lossy one. Two intervals can report the same 6 percent and behave nothing alike: one a handful of large, near-circular, possibly touching vugs that form a flow path, the other a scatter of tiny isolated pores that store fluid but never connect. The bulk percentage cannot distinguish them, because by the time you have a percentage you have already thrown away every individual vug, its size, its shape, and its orientation.

This whitepaper is about putting that information back. In our work with a mid-sized Middle East carbonate operator, we built an automated computer-vision pipeline that does not produce a percentage; it produces a catalogue — every vug detected as its own object, with its own area, its own circularity, and its own azimuth — and then lets the statistics fall out of that catalogue rather than standing in for it. It is written for the petrophysicist and reservoir-quality specialist who consume the interpretation, the core analysts who calibrate it, and the secondary-porosity leads deciding whether per-vug quantification is worth the engineering.

Why "not deep learning" was the right engineering call

It would be fashionable to reach for a segmentation network here. We deliberately did not, and the reason is an engineering one rather than a doctrinal one. A supervised deep-learning vug segmenter needs a corpus of labelled vug masks to learn from, and at the start of this work that corpus did not exist — the only ground truth was the operator's hand-marked interpretation-software percentages, which are bulk, not per-vug. Training a black-box model against a bulk label to recover per-vug structure is the wrong tool for the wrong target.

A classical pipeline inverts the dependency. Every stage is an interpretable operator with a parameter a geoscientist can reason about — a threshold, a block size, a circularity ratio — so the method is auditable, tunable on a handful of wells rather than hundreds, and defensible in a review where "the network said so" is not an acceptable answer. And it produces exactly the artefact a future deep-learning effort would need: a per-vug mask catalogue that can serve as ground truth for a supervised or semi-supervised segmenter later. The classical pipeline is not a competitor to the deep-learning roadmap; it is the bootstrap that makes the roadmap possible.

THE INFORMATION A PERCENTAGE DISCARDS

An interpretation-software vug percentage is a scalar per interval. The pipeline in this whitepaper emits, per 0.1-metre interval, the total vug count, total and mean vug area, the standard deviation of area, and the full area, circularity, and azimuth spectra. The same interval that was one number becomes a small distribution — and reservoir quality lives in the distribution, not the number.

The pipeline, stage by stage


The system is a single-pass image-processing chain. Input is a static image log — we use static rather than dynamic deliberately, because dynamic normalisation enhances the bedding planes and fractures we are trying to *exclude*, whereas static preserves the contrast a vug presents against its matrix. A single image log is often longer than two hundred metres; it is rescaled to a 0–255 intensity range, segmented into one-metre processing zones, and quantified at every 0.1-metre interval. Image logs cover only about 80 percent of the borehole wall, so the unmeasured pixels arrive coded as a `-9999` sentinel and are carried through as NaN rather than silently treated as black.

Local-variation enhancement by top-K mode subtraction. A raw image log has slowly varying background trends that defeat a global threshold. We remove them with a mode-based local enhancement: take the K most frequent intensity modes of a patch and subtract them, leaving the local departures — the candidate vugs — standing out from a flattened background. We use K = 5 top modes with an inter-mode spacing of 5 intensity levels. That

spacing is not arbitrary: at a spacing of 1, consecutive mode-subtracted images are over 90 percent duplicate content; at a spacing of 5 they isolate genuinely distinct features, which is what justifies the choice.

Gaussian-modulated adaptive thresholding. With the background flattened, each pixel is thresholded against a Gaussian-weighted local mean rather than a single global cut. Two parameters govern it: a block size — the neighbourhood over which the local threshold is computed — and a constant offset `C` subtracted from that local mean. We tested block sizes of 11, 21, 31, 41, and 51 pixels and settled on 31 as the default, which at this image resolution covers about 46.58 square centimetres of borehole wall; 11 pixels covers only 5.87 square centimetres and fragments larger vugs, while 51 covers 126.08 and over-merges them. The offset `C` was swept from -2 to +2 and set to the patch's mean intensity (equivalently, -1) as the default.

Contour extraction and the geometric gate. The thresholded binary image is traced into closed contours with the Suzuki-Abe border-following algorithm. Each contour's area is computed by the shoelace formula (Green's theorem on the boundary) and its circularity by comparison against its minimum enclosing circle. This is the first place geoscience enters as a constraint: a vug is a roughly rounded cavity, so we gate contours by a circularity ratio between 0.3 and 1, discarding the long, thin, low-circularity contours that a bedding plane or fracture produces. Area and circularity together form the geometric gate — the structural filter that says "this shape could be a vug" before any contrast check runs.



The pipeline's single geometric decision, made tactile. Each contour the detector traces sits on a circularity axis — 0.28 (elongated, dissolution-aligned) to 0.85 (near-circular). A circularity gate of 0.3–1.0 rejects linear fractures and wellbore-parallel artifacts while keeping true vugs. Drag the gate: contours left of it fall out as rejected fractures (grey), contours right of it enter the vug catalog (teal), and the three orange contours are vugs that meet every geometric criterion the expert applied — yet manual picking missed across two of the validation wells. The gate bounds, circularity span, 1–12 cm² area, and three recovered intervals are the article's own; each contour's exact coordinate is schematic (a plausible population, not a published catalog).

The catalogue above is the heart of the method made interactive. Each contour is a point in area-versus-circularity space; sliding the gate shows how the geometric constraint admits the rounded, vug-shaped detections and rejects the elongated artefacts — and, crucially, how a gate set too aggressively starts discarding real dissolution vugs along the low-circularity edge. The gate is a dial a geoscientist sets, not a constant the model hides.

Aggregation: de-duplicating across modes. Because we run the threshold across several mode-subtracted images, the same vug is often detected more than once. We aggregate by two checks: a centroid-difference threshold (contours whose centres lie within $\tau = 5$ pixels are treated as the same object) and an intersection-over-union merge at a 20 percent overlap threshold. The centroid spacing was itself tuned — a contrast spacing of 1 removed


roughly 250 redundant contours, a spacing of 2 a further 120, and by a spacing of 5 the count had stabilised at around 30 removals, the point of diminishing returns where distinct vugs are no longer being merged.

The false-positive problem, and the two-stage filter that solves it

Everything to this point produces candidate vugs. It also produces false positives, and in a carbonate image log the false positives are not random noise — they are *other geology*. Conductive bedding planes and fractures present as dark features too, and a geometric gate alone will let through the ones that happen to look rounded in a single patch. The two filters that suppress them are where the pipeline earns its keep, and both are deliberately physics-aware.

Laplacian-variance contrast filtering. A true vug is a coherent dark blob with a sharp boundary against bright matrix; a smeared artefact along a bedding plane is not. The Laplacian variance inside an expanded bounding box around each contour measures exactly that boundary sharpness. We expand the bounding box by 10 percent on each side — 20 percent total — and reject contours whose Laplacian-variance contrast falls below a threshold. That threshold was swept from 0.05 to 0.75 and set to 0.25 by default; it is one of the two parameters we expect to re-tune per well.

Mean-intensity differential filtering. The second filter is the one that makes the method portable across mud systems. A real vug should be darker — more conductive — than an annulus of rock immediately around it; a feature that is no darker than its surround is not a vug. We compare the mean intensity inside the contour against the mean intensity of an enlarged circle of radius $2R$ around it, and reject contours that fail the contrast. The sign of that comparison depends on the mud: in a water-based-mud well the conductive vug is dark and the inner mean is lower, and the filter checks accordingly. The mean-based threshold was swept from 0.1 to 2.0 and set to 0.8 by default. This is the filter that physically confirms a candidate is a vug — and as a quality check, the detected vugs do shift the local resistivity reading toward lower values, exactly as a genuine conductive cavity should.



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.

The recovery curve above is the operational point of this stage. On clean intervals the geometric gate is nearly sufficient and the false-positive filters barely fire. On conductive zones and bedding-plane-rich sections, the same filters do heavy lifting — suppressing the features that would otherwise inflate the vug count — and the curve shows how detection quality is recovered as the contrast-based suppression engages on exactly the intervals where naive thresholding fails.

One parameter set, three wells, two tools

The test of an interpretable pipeline is not a single-well demo; it is portability. We validated on three wells chosen to stress that portability: a vertical well logged with a high-resolution borehole image log, a horizontal well logged with a compact microresistivity tool roughly ten kilometres away, and a second vertical well, also imaged with a high-resolution

borehole image log, roughly twelve kilometres away. The point of the horizontal well is that it changes almost everything an image-processing pipeline is sensitive to — tool response, borehole geometry, and the orientation of the vugs in the unrolled image — and the point of the spread is that these are genuinely different rock and acquisition conditions, not three intervals of one log.

The universal parameters held across all three wells: the number of modes, the inter-mode spacing, the centroid-difference threshold, the IoU threshold, the Laplacian bounding-box expansion, and the mean-based radius. Only the two contrast thresholds and the block size needed re-tuning. The horizontal well ran with a Laplacian threshold of 0.7, a mean-based threshold of 1.0, and an offset `c` of -1.0; the second vertical well ran with a Laplacian threshold of 0.15, a mean-based threshold of 0.4, and a block size widened to 51 pixels. That is the whole adaptation budget — two thresholds and a window size — to move a carbonate vug detector from a vertical high-resolution borehole image log to a horizontal compact-microresistivity log. A model that needed full retraining to cross that gap would not be a tool; it would be a project.

WHAT PORTABILITY BUYS THE OPERATOR


The adaptation cost between a vertical high-resolution-image-log well and a horizontal compact-microresistivity well was two contrast thresholds and one block size. Everything geological in the pipeline — the mode subtraction, the circularity gate, the aggregation logic — transferred unchanged. That is the difference between a method an operator's own petrophysicists can run on the next well and a method that requires the vendor back on site.

What the catalogue reveals that a percentage cannot

Run the pipeline across the validation wells and the per-vug statistics tell a story the bulk percentage never could. Individual vug areas span roughly 1 to 12 square centimetres, but they are not uniform across that range — they concentrate between 1 and 3.5 square centimetres, a right-skewed distribution of many small vugs and a thin tail of large ones, with the per-interval spectra typically sitting between 1 and 6 square centimetres. Circularity runs from about 0.28 to 0.85 and peaks between 0.45 and 0.7 — semi-circular

dominant, consistent with dissolution vugs that are rounded but not perfectly so. And the azimuth spectrum, binned into the four quadrants of the borehole, reveals whether the vugs cluster on one side of the wellbore or distribute evenly — directional information a single percentage erases entirely.

This is the argument for the whole approach in one paragraph. A reservoir-quality specialist looking at "6 percent vug porosity" learns one thing. A specialist looking at "6 percent, made of forty-two vugs, mean area 2.1 square centimetres, right-skewed, circularity peaking at 0.55, azimuth clustered in the upper-left quadrant" learns whether those vugs are likely to connect, whether they are dissolution-driven or fabric-selective, and where in the wellbore to look. The percentage is a summary of the catalogue; the catalogue is the thing worth having.



The *whitepaper's opening thesis, made tactile*: a bulk vug percentage is a lossy scalar. Drag the texture lever and the bulk vug percentage at the top does not move a decimal, while the three statistics the pipeline actually emits per 10 cm interval, vug count, mean per-vug area, and standard deviation of area, swing across the whole range. At the few-large end, a handful of near-circular vugs carry the porosity and are likeliest to touch and form a flow path; at the many-small end, the same porosity is spread across a right-skewed swarm of isolated pores that store fluid but rarely connect. The per-vug area histogram on the right redistributes accordingly across the observed 1 to 12 square-centimetre span, concentrated 1 to 3.5. The single orange element is the flow-connectivity verdict, the one thing that argues, and it is a qualitative index of the article's claim rather than a measured number. Sourced from the engagement archive: the definition of bulk vug percentage as vug area over interval area, the 1 to 12 square-centimetre observed range concentrated 1 to 3.5 and right-skewed, and the per-10-centimetre emitted statistics; the fixed 6 percent value, the two archetype vug counts, and the exact per-vug areas placed in the histogram are illustrative, a plausible decomposition rather than a published catalogue.

The lens above holds the bulk percentage fixed and lets you drag the fabric underneath it. The scalar never moves; the count, the mean vug area, and the standard deviation of area swing across their whole range, and the per-vug area histogram redistributes from a thin tail of large, connectable vugs to a right-skewed swarm of small isolated ones. That is the lossy step made visible: two intervals that report the same percentage, and flow nothing alike, are indistinguishable the moment you collapse the distribution into the number.

Catching what the interpreter missed

The most consequential validation result is also the simplest to state. At several intervals across the validation wells, the pipeline detected vugs that the manual interpretation-software ground truth had not marked. These were not spurious — they passed the circularity gate and both contrast filters, and the local resistivity confirmed them — they were genuine dissolution vugs that a human picking by eye across a two-hundred-metre log had simply missed. Manual interpretation is careful work, but it is also fatiguing, sequential, and inconsistent across a long log and across interpreters. An automated detector applies the same gate to every centimetre with the same patience at metre two hundred as at metre two, and that consistency is where the additional vugs come from.

It is worth being precise about what this does and does not claim. The pipeline is not "more accurate than the expert" in some abstract sense; the expert remains the ground truth and the arbiter, and the system also produced the occasional false positive — a fracture near a vuggy zone briefly mis-identified — that an expert correctly rejected. What the pipeline provides is *exhaustive, uniform first-pass detection* that the expert then reviews and signs off, rather than asking the expert to do the exhaustive detection by hand. The expert is retained for judgement; the machine is assigned the patience.

THE RIGHT DIVISION OF LABOUR

This is not an autonomous interpreter and it should not be sold as one. It is a detection-and-quantification engine whose output a petrophysicist reviews. The value is not that it removes the expert; it is that it changes what the expert spends time on — from marking every vug by hand to adjudicating a uniformly-detected catalogue and the handful of edge cases it flags.

The engineering around the algorithm

A pipeline that works in a notebook on one interval is a result; a pipeline that runs on a two-hundred-metre log unattended is a product, and the gap between them is software engineering. Three pieces of that engineering matter here.

First, **the NaN discipline**. Because image logs cover only about 80 percent of the borehole and the gaps arrive as a `-9999` sentinel, every operator in the chain — the mode subtraction, the thresholding, the area integration — has to treat NaN as NaN and not as a dark pixel that looks like a vug. Getting that wrong inflates the vug count along every gap in the log. It is unglamorous and it is load-bearing.

Second, **the speed envelope**. The prior published approach to automated vug detection used a path-morphology method that ran on the order of five minutes per metre — too slow to deploy across a multi-well programme. A chain of vectorised image operations runs orders of magnitude faster, which is what turns "a method" into "a method you can run on every well." We are deliberately not publishing a single headline seconds-per-metre figure here, because an earlier draft's speed claim was tightened in peer review and we would rather under-claim than overstate; the operationally relevant fact is that the pipeline is fast enough to process a field's worth of logs, where the path-morphology baseline was not.

Third, **reproducibility by construction**. Every parameter that governs a run — the K modes, the block size, the offset, the circularity ratio, the two contrast thresholds, the per-well overrides — is an explicit, recorded constant, not a buried magic number. A run is fully described by that parameter set plus the input log version, which means a result can be reproduced exactly and a parameter change can be attributed to its effect. That is the same discipline that lets the per-well overrides above be stated in a sentence: they are values in a config, not lore in a notebook.

What this is the foundation for

The deepest reason to build the classical pipeline first is that it manufactures the asset the next phase needs. A supervised deep-learning vug segmenter is a reasonable long-term goal — it would absorb the per-well tuning into learned weights and could capture vug morphology a fixed circularity gate approximates. But it needs labelled per-vug masks to train on, in volume, and those masks are exactly what this pipeline produces: a reviewed, geoscientist-signed-off catalogue of vug contours across multiple wells and two tool types. The interpretable method is the data engine for the learned method. Run it across a field, have the petrophysicists adjudicate its output, and the by-product is a training corpus that did not exist before — which is how a programme moves from classical computer vision to deep learning without betting the first phase on labels it does not yet have.

That is the through-line of this work and of the wider programme it belonged to: treat the geoscience as the specification, the algorithm as its reproducible enforcement, and the production layer as what makes it survive contact with a real field. The vug pipeline is a small, sharp instance of that philosophy — interpretable where interpretability matters, fast where speed gates deployment, and pointed deliberately at producing the ground truth its own successor will learn from.

WHAT THIS WHITEPAPER ARGUES

- 1 A bulk vug percentage is a lossy scalar: two intervals at the same percentage can flow very differently. Per-vug quantification recovers the area, circularity, and azimuth distributions that actually control reservoir quality.
- 2 The pipeline is classical CV by design, not dogma: mode subtraction → Gaussian adaptive thresholding → Suzuki-Abe contours → area+circularity gate → centroid/loU aggregation → Laplacian + mean-intensity false-positive filters. Every stage is an interpretable, tunable operator.
- 3 The two-stage false-positive filter is physics-aware: Laplacian-variance contrast rejects smeared bedding-plane artefacts, and mean-intensity differencing (mud-sign-aware) confirms a candidate is genuinely more conductive than its surround — suppressing the conductive features a naive thresholder mistakes for vugs.
- 4 Validated on three wells spanning vertical and horizontal high-resolution borehole image logs acquired with two different microresistivity imaging tools, using one shared parameter set — only two contrast thresholds and a block size re-tuned per well — and it recovered dissolution vugs the manual interpreter missed.
- 5 The interpretable pipeline is the bootstrap for deep learning: its reviewed per-vug catalogue is exactly the labelled-mask ground truth a future supervised vug segmenter needs and that bulk interpretation-software percentages cannot provide.

References

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 underlying the detection stage. [https://doi.org/10.1016/0734-189X\(85\)90016-7](https://doi.org/10.1016/0734-189X(85)90016-7)

Li, et al. Path-Morphology Approaches to Borehole-Image Feature Extraction (2019). The prior automated approach whose per-metre cost motivated a faster vectorised pipeline.
<https://doi.org/10.1190/segam2019-3216905.1>

Lucia, F. J. Carbonate Reservoir Characterization: An Integrated Approach. Springer (2007).
Reference on vuggy secondary porosity and its control on carbonate reservoir quality.
<https://doi.org/10.1007/978-3-540-72742-2>

Beyond Bulk Percentages: AI-Assisted Individual Vug Quantification in Carbonate Reservoirs

Published August 2024. Generated 2026-07-22.

<https://www.earthscan.io/whitepapers/ai-vug-quantification-carbonate-beyond-wellcad>

Copyright EarthScan. All rights reserved.