

Anatomy of a Top-K Mode-Subtraction Vug Detector: Five Enhanced Copies of One Log

Automated vug quantification on carbonate borehole-image logs has a published body of work on each of its stages in isolation: mode subtraction to suppress the matrix, adaptive thresholding to resolve pores, a circularity gate to reject fractures, a false-positive filter to clean the catalog. What a reviewer asks for is different: not the best stage, but the whole pipeline stated precisely enough to reproduce and defend. This whitepaper is that statement. One borehole-image patch is min-max scaled to 0-255, then turned into five mode-subtracted copies (top-K modes $K=1$ to 5) at a minimum intensity spacing of $\Delta m = 5$ so the copies are genuinely distinct. Each copy is adaptively thresholded with a 31×31 Gaussian-modulated block, a window that covers 46.58 cm^2 of borehole wall, at a one-pixel step. Contours are extracted from every copy and each candidate is validated with a Laplacian-variance and mean-difference test on a bounding box expanded 10 percent per side, with a mean-based enlarged radius of $2R$ and cross-copy duplicates merged at an intersection-over-union of 0.2. The output is an interval signature: per-vug area, predominantly 1 to 6 cm^2 ; circularity, distributed 0.28 to 0.85 with a semi-circular peak in 0.45 to 0.7; and azimuth, resolved into four quadrants. We present the circularity ridgeline spectrum and the estimated-vs-manual vug track as the two instruments that make the case, and we keep to the constants that survived review: $k=5$, $\Delta m = 5$, block 31, IoU 0.2, 10-percent expansion, runtime near 2 minutes per 2 m interval with 100 threshold values for path opening. The argument is that the pipeline's defensibility comes from the validation stage, not the detection stage.

Tannistha Maiti

Senior AI Researcher

Narendra Patwardhan

Research Collaborator

VUG-DETECTION / MODE-SUBTRACTION / ADAPTIVE-THRESHOLDING / LAPLACIAN-VARIANCE / CONTOUR-VALIDATION / CIRCULARITY

CONTENTS

01	One log, five copies: the shape of the pipeline
02	The threshold has a physical size
03	Contours are candidates, not conclusions
04	The validation stage is where the pipeline earns its number
05	Reading the output: circularity is the signature, not the contour
06	Azimuth: the direction the pores face
07	The estimate against the expert
08	What survived review, and why it is stated this way
09	Limitations
10	References

A vug detector is easy to demo and hard to defend. The demo is a dark blob on a borehole-image patch, a threshold, a green outline drawn around the blob, and a number: this interval is 2 percent vug. It looks finished. Then a reviewer reads it, and the questions that come back are not about the blob. They are about the four decisions no one wrote down: which threshold, computed over what window, applied to which version of the image, and how you know the green outline is a dissolution pore rather than a shadow at the edge of a fracture. The pipeline that answers those questions is a different object from the one that produced the demo. It has stages the demo skipped, constants that had to be pinned rather than tuned per well, and a validation step that exists only because a reviewer refused to take the detection step on faith.

This whitepaper is the vug detector as it survived that process. It is written for the petrophysicist who has to sign the porosity number, the chief geoscientist deciding whether a classical computer-vision pipeline belongs in an interpretation workflow, and the data scientist who will maintain it. We built and validated it in an engagement with a mid-sized Middle East carbonate operator, and the numbers in it are the ones that made it into the final, reviewer-accepted manuscript. The individual stages have been described before, in our own earlier notes and in the wider literature; where a mechanism is already published we compress it to a pointer and spend the words on how the stages fit together and, above all, on the validation stage that the earlier stage-by-stage treatments did not have to defend as a whole.

The framing to hold onto through what follows is a distinction between two kinds of work the pipeline does. The first kind is detection: finding dark regions on the borehole image. The second is validation: deciding which dark regions are dissolution pores. The published literature, and the demo, are mostly about the first kind, because finding dark regions is the visible part and the part that produces a satisfying green outline. The reviewers cared about the second kind, because the second kind is where a vug count becomes a defensible measurement rather than a blob tally. The stages before validation are stated here for completeness and to fix the constants; the weight of the paper sits on the validation stage and on the two instruments that read its output against the physics and against the expert.

One log, five copies: the shape of the pipeline

The pipeline takes one thing in and puts one thing out. In goes a patch of a static borehole-image log, roughly a metre of borehole wall unrolled into a rectangle of intensities. Out comes a signature per interval: how much vug, of what size, how round, facing which way. Everything between is deterministic. Given the same patch and the same constants the pipeline returns the same contours every time, which is the property that lets a petrophysicist defend the number a year later.

The first stage is unglamorous and load-bearing. The raw intensities of a borehole-image patch do not live on a fixed scale; static image ranges drift with tool, gain, and formation, and a compact microresistivity tool in a horizontal well can produce a static range wildly different from a high-resolution imager in a vertical one. So the patch is min-max scaled to the standard 0 to 255 range before anything else touches it. This is not cosmetic. Every downstream constant, from the mode spacing to the threshold offset, is expressed on that 0 to 255 scale, so the rescale is what lets a single set of constants apply across wells and tools rather than being re-tuned each time.

Then the log is copied. Not duplicated, enhanced: the pipeline generates five mode-subtracted copies of the scaled patch, the top-K modes with K running 1 to 5. The mechanism is the one we described in the earlier technique note (see the pointer below): the carbonate matrix is the majority class in the intensity histogram, so subtracting the most frequent intensity modes deletes the background by its own statistics and leaves the darker, rarer vug pixels in relief. What matters for the pipeline anatomy is the number of copies and the spacing between them. Five copies, at a minimum intensity-level separation of $\Delta m = 5$, so the copies are genuinely distinct slices of the intensity range rather than five near-identical images. Take the top five modes literally, with no spacing, and the copies collapse onto each other and the contour stage downstream re-traces the same pores five times. Force $\Delta m = 5$ and each copy surfaces a complementary slice of the vug population. That is the whole reason the detector is five copies of one log and not one copy: coverage of the pore population without duplication, bought with two integers.

For the mechanics of why mode subtraction beats a global threshold, and the histogram intuition behind it, see our earlier deep-dive [Top-K Mode Subtraction: Isolating Dissolution Pores from Carbonate Background](#). This paper takes that mechanism as given and asks what has to happen to its output before a reviewer will accept the vug count it produces.

The threshold has a physical size

Each of the five copies is adaptively thresholded. Adaptive, because a single global cutoff cannot separate vug from matrix across a whole well when the matrix intensity itself drifts with depth and lithology. The threshold for each pixel is computed from a Gaussian-modulated weighting of its neighbourhood, a block, with a small offset constant, and the block is where the geology enters.

Block size was swept over odd values and fixed at 31 by 31 pixels. The number is not arbitrary and it is not unitless. A 31-pixel block covers about 46.58 cm² of borehole wall, comfortably larger than a typical vug but small enough to track the matrix's local variation. Sweep down to an 11-pixel block, covering 5.87 cm², and the window is too small to establish a stable local background and fragments a single vug into pieces; sweep up to 51 pixels, covering 126.08 cm², and the window is large enough to smear several nearby vugs into one threshold decision. The block is sized deliberately to be bigger than the signal and smaller than the trend, and the step is one pixel, so the window slides across the whole patch without gaps. The offset constant defaults to the mean intensity of the current patch, which keeps the threshold anchored to the local level rather than a fixed absolute value.

The eight interpretable constants that govern this stage, and the argument that their interpretability is exactly what makes the detector a better porosity-QC tool than a learned mask, are laid out in our [eight-parameter adaptive-threshold note](#); we do not re-derive them here. The stage-flow instrument below shows where thresholding sits in the full pipeline and what it hands forward.



The vug detector as it survived peer review, drawn as the six stages it runs on one borehole-image patch: min-max scale to 0-255, generate k mode-subtracted copies at delta-m 5 (top-K modes, K=1 to 5), 31x31 Gaussian-modulated adaptive thresholding (a block-31 window covers 46.58 cm2 of borehole wall), extract contours from each copy, validate every candidate with Laplacian-variance and mean-difference filtering on a 10-percent-expanded bounding box (mean-based radius 2R, cross-copy duplicates merged at IoU 0.2), and emit the output signature: per-vug area, circularity, and azimuth. The copies lever (drag k from 1 to 5) fans out the mode-subtracted copies at stage 2; because delta-m 5 keeps those copies genuinely distinct, each surfaces a complementary slice of the vug population, and the validation stage merges what overlaps so the whole well still resolves to one signature per interval. The orange panel on the right is the only element that argues: the area, circularity, and azimuth the whole pipeline exists to produce. The stage numerics are sourced from the reviewer-response archive; the redundant-merged fraction on the lever is an illustrative reading of the duplicate-merge behaviour.

The stage flow makes the pipeline's shape concrete: one scaled patch, five mode-subtracted copies fanning out at the second stage, an adaptive threshold on each, contours extracted from each, and then the stage that the demo never had, validation, before the output signature on the right. Drag the copies lever and the fan widens or narrows; the point the flow argues is that widening it does not multiply the output, because the validation stage merges what the extra copies duplicate. The number of copies is a coverage decision, not an output-count decision. That distinction is the difference between a detector that over-reports vugs in proportion to how many copies you happen to run and one that reports the same population regardless.

Contours are candidates, not conclusions

Thresholding turns each copy into a binary image, and contour extraction turns each binary image into a list of closed outlines, one per connected dark region. In the demo, this list is the answer. In the pipeline that survived review, it is a list of candidates, because a dark contour is not yet a vug. Three things produce dark contours: dissolution pores, which are what we want; fractures and bedding traces, which are elongate and cut across the image; and tool artefacts, which run parallel to the wellbore. The contour stage cannot tell them apart. It just finds closed dark regions on five separate copies, which means the same physical pore often appears as a contour on more than one copy, and every one of those contours is, at this point, unvalidated.

So the pipeline has, coming out of contour extraction, a redundant and impure catalog: the real pores, several times over, mixed with fractures and artefacts. Everything published on this method up to the detection stage stops roughly here and reports a vug count. The reviewers did not accept that, and they were right not to. A count taken off the raw contour list is a count of dark blobs, not of dissolution pores, and it double-counts the ones that survived on multiple copies. The rest of this paper is about the stage that turns the candidate list into a catalog you can defend.

The validation stage is where the pipeline earns its number

Each candidate contour is validated before it enters the catalog, and the validation is two physical tests run on a deliberately enlarged view of the contour, not on the contour itself.

The enlargement comes first. The bounding box of each candidate contour is expanded by 10 percent on each side, 20 percent total, and a mean-based enlarged radius of $2R$ is used so the test region is a genuine neighbourhood around the candidate rather than just its tight outline. Concretely, a candidate with a characteristic radius R is tested over a region of radius $2R$ centred on it, so the surround the tests read is as large as the candidate itself:

VALIDATION REGION AROUND A CANDIDATE OF RADIUS R

Formula

LaTeX

$$R_{\text{test}} = 2R, \quad \text{bbox}_{\text{expanded}} = (1 + 2 \times 0.10) \text{bbox}$$

Copy LaTeX

This matters because the discriminating information about whether a dark region is a real pore lives at its margins and in its immediate surroundings, not inside it. A pore has a sharp, textured boundary against the matrix; a smear or a shadow does not. Testing only the interior of a tight contour throws that information away, which is why a tight-contour count over-reports: it never looks at the one place, the rim, where a false positive gives itself away.


On that expanded box, two tests run. The first is a Laplacian-variance test. The Laplacian responds to second-order intensity change, edges and texture, and its variance over the expanded box measures how much genuine structure sits at the candidate's boundary. A real pore, with a crisp rim against the carbonate, produces high Laplacian variance; a low-contrast smear or a gradual intensity drift produces low variance and is rejected. The second is a mean-difference test: the mean intensity inside the candidate is compared against the mean of the expanded surround, and a real pore, being a genuine low-resistivity cavity, shows a real intensity contrast against its neighbourhood, while an artefact riding on the local background shows little. A candidate has to pass both to survive.

Then the redundancy is resolved. Because five copies were thresholded independently, the same pore appears as several candidates; these are merged when their bounding boxes overlap past an intersection-over-union of 0.2. The merge is what makes the copy count a coverage decision rather than an inflation of the count: run more copies and you find each pore on more of them, but the IoU-0.2 merge collapses those hits back to one entry. The full contrastive and vicinity logic of this false-positive stage, including the circularity-and-area gate that runs alongside the Laplacian and mean-difference tests, is documented in our [false-positive filtering case study](#); the anatomy point here is narrower and, we think, the most important thing in the pipeline. The detection stage decides what is dark. The validation stage decides what is a vug. The reviewers' confidence, and ours, rests on the second stage, and the second stage is the one the demo did not have.

Reading the output: circularity is the signature, not the contour

The validated catalog is counted every 10 cm down the well, and the count is not a single number but a signature: per-vug area, a circularity spectrum, and an azimuth spectrum, per interval. Of the three, circularity is the one that carries the most argument, because it is the geometry that separates a dissolution pore from everything else that produced a dark contour.

Vugs are roughly equant, semi-circular to circular, because dissolution eats rock in all directions from a nucleation point. Fractures are elongate. So circularity, the ratio that runs from near 0 for a line to 1 for a perfect circle, is the physical invariant that says vug. Across the validation wells the surviving vug population showed a circularity distribution spanning 0.28 to 0.85 with a clear peak in the 0.45 to 0.7 band, semi-circular dominant, exactly the diagenetic signature you expect. The important move is that this is read as a spectrum per interval, not as a per-contour label. No single contour's circularity proves it is a vug; the shape of the distribution across an interval is what does. An interval whose circularity mass sits in 0.45 to 0.7 is an interval of dissolution porosity. An interval whose mass sits to the left, near 0.3, is an interval of fractures that leaked through the earlier gates, and it reads that way at a glance.



The circularity spectrum the pipeline emits every 10 cm, stacked down a short depth window as a ridgeline. Each ridge is the circularity histogram of the vugs validated in that interval; the whole stack argues that real dissolution pores pile up in a semi-circular band, so a circularity spectrum with its mass in 0.45-0.7 is the signature of vugs rather than of elongate fractures, which fall to the left. The observed distribution spans 0.28 to 0.85 with the peak in 0.45-0.7. The orange band marks that semi-circular peak across every ridge, and the gate lever sweeps the circularity accept threshold: keep it left of 0.45 and the whole peak survives, push it into the band and you start discarding real vugs. The distribution bounds and the peak band are sourced from the reviewer-response archive; the per-interval ridge shapes are illustrative renderings of that sourced distribution.

The ridgeline stacks the circularity spectrum of successive 10 cm intervals down a short window, and the orange band marks the 0.45 to 0.7 semi-circular peak across all of them. The instrument argues the claim directly: the ridges pile up in the band, interval after interval, which is what a genuine vug population looks like. The gate lever sweeps the circularity accept threshold; keep it left of 0.45 and the whole peak survives, push it into the band and you begin discarding real pores. That is the tension the spectrum makes visible. A gate tuned too aggressively to kill fracture leakage will also amputate the low-circularity tail of the real vug population, and the ridgeline is how you see the cost before you pay it. Circularity is not a filter you set once and forget; it is a distribution you read.

The area channel completes the size half of the signature. Vug areas came out predominantly in the 1 to 6 cm² range, which is consistent with the block sizing decision: a 31-pixel block covering 46.58 cm² is comfortably larger than the pores it has to resolve, so the threshold sees each pore against a stable local background rather than the pore filling the window. The area distribution is a check on the whole chain. If the reported areas had clustered near the block area, it would mean the threshold was resolving the window rather than the pore, and the block would be wrong. They did not; they sat well below it, with the bulk between roughly one and six square centimetres against a window of almost fifty.

That internal consistency is the kind of check the validation stage buys and the raw contour count cannot offer. Area, circularity, and block size are three numbers that have to agree for the pipeline to be measuring what it claims: the block must be larger than the pores, the pores must be roughly equant, and the areas must sit below the block. When all three hold, the signature is coherent; when one drifts, the coherence breaks and the drift is visible in the distributions rather than buried in a scalar. A petrophysicist reading the area spectrum, the circularity ridgeline, and the block area together is reading a self-checking measurement, which is a stronger thing to sign than a single porosity percentage with no way to interrogate how it was formed.

Azimuth: the direction the pores face


The third channel of the signature is the azimuth spectrum, and it is the one that connects the vug catalog back to the structural picture. Every contour on an unrolled borehole image has an angular position, because the horizontal axis of the image is azimuth around the

borehole. So a vug catalog carries, for free, the compass direction each pore faces, and aggregating those directions across an interval produces an azimuth spectrum split into four quadrants: 0 to 90, 90 to 180, 180 to 270, and 270 to 360 degrees.

The azimuth spectrum is diagnostic in a way the area and circularity channels are not. Area and circularity say how much vug and what kind. Azimuth says whether the dissolution has a preferred orientation, which is a question about the diagenetic and structural history of the interval. A vug population smeared evenly across all four quadrants reads as fabric-selective dissolution with no directional control. A population that clusters in one or two quadrants suggests dissolution guided by a fracture set or a bedding-parallel fabric, which is exactly the kind of coupling a geomodeller wants flagged. The pipeline does not interpret this; it reports the spectrum and leaves the interpretation to the geoscientist. But reporting it at all is only possible because the earlier stages produced a validated catalog of individual pores with positions, rather than an interval-averaged porosity number. The signature is richer than a scalar because the pipeline counted objects, not pixels.

The estimate against the expert

None of the anatomy matters if the numbers the pipeline emits do not track what an interpreter would pick by hand. The test that mattered to the reviewers, and to us, was per-interval agreement: take the estimated vug percentage the pipeline produces for each short sub-interval and set it beside the manual, ground-truth vug percentage the expert picked, interval by interval, and look at where they agree and where they part.



A depth track that reads the pipeline against the expert. For each short sub-interval the estimated vug percentage the mode-subtraction detector emits (filled teal) is plotted beside the manual, ground-truth vug percentage the interpreter picked (hollow marker), joined by a connector whose length is the disagreement. The pairs are sourced: 1.011 vs 1.064, 1.246 vs 1.071, 0.869 vs 2.815, 2.354 vs 2.008, 0.940 vs 1.630, 0.584 vs 0.669, 1.049 vs 0.297 percent. The tolerance lever sweeps an agreement band; connectors that exceed it turn orange, the one element that argues, marking the intervals a reviewer actually reads. Most intervals sit inside a narrow band, and the departures point in a readable direction: in side-by-side review some were real dissolution pores the manual pass had left flat. The paired percentages and the 1-to-6 cm² area range are sourced from the reviewer-response archive; the depth offsets are anonymised relative labels.

The track reads the two side by side. For most intervals the estimate sits close to the manual pick, inside a narrow band, which is the result you need before the signature is usable as anything but a screen. The departures are the interesting part, and the tolerance lever is there to isolate them: widen the band and the flagged intervals fall away, narrow it and the disagreements light up in orange. Two things about those departures matter. First, they are a minority; the bulk of the track is agreement. Second, and this is the finding that changed how we present the tool, some of the departures were not the pipeline being wrong. In side-by-side review, intervals where the estimate read vug and the manual pick read flat turned out, on a second look, to contain real dissolution pores the manual pass had missed. The pipeline was catching porosity, not inventing it. That does not make every departure a save; some are genuine over-reads at the margins. But it reframes the

disagreement column from a list of errors into a list of intervals worth a second human look, which is precisely the right role for an automated first pass in a workflow where the expert stays in the loop.

That reframing is the argument for shipping this as a screen-and-review tool rather than an unattended oracle. The pipeline does the mechanical work across kilometres of log, produces a defensible signature per interval, and flags the handful of intervals where its number and the expert's would part. The expert spends time on those, and on sign-off, instead of on picking every pore by hand. The runtime supports this posture: about 2 minutes per 2 m interval, with 100 threshold values swept for the path-opening step, which is fast enough to run the whole well and slow enough only where the interval genuinely needs the sweep.

What survived review, and why it is stated this way

The constants in this paper are the ones that made it through peer review unchanged, and it is worth naming why they are stated as fixed numbers rather than tuning ranges. $k=5$ copies, delta-m 5 spacing, block 31 by 31, threshold offset at the patch mean, bounding-box expansion 10 percent per side, mean-based radius $2R$, IoU merge 0.2. These were held universal across wells and across two different imaging tools, a vertical well logged with a high-resolution imager and a horizontal well logged with a compact microresistivity tool kilometres away. A method whose behaviour is governed by a handful of interpretable constants that do not move well to well is a method a petrophysicist can specify once and reason about, which is what the reviewers were really testing when they asked for the pipeline rather than the demo.

The deeper reason to state it this way is that the defensibility of the vug number does not come from the detection stage at all. Mode subtraction and adaptive thresholding are good at finding dark regions, but finding dark regions is the easy, already-published part. The number becomes defensible at the validation stage, where a Laplacian-variance test and a mean-difference test on an expanded neighbourhood decide, per candidate, whether a dark region has the physical signature of a dissolution pore, and where the IoU merge collapses the redundancy that five copies necessarily produce. A reviewer who accepts the vug count is accepting that stage. The comparison to a purely morphological baseline makes the point from the other direction: an approach that leans entirely on morphology to

characterise vugs pays for it in runtime and in false positives it has no principled way to reject, whereas the validation stage here rejects them on a physical test and does so fast enough to run at scale.

We are not claiming this is the only way, or that a learned segmentation model could not eventually do the same job. We are claiming that for this problem, on this data, with the labels that existed, the pipeline that survived review earns its number at the validation stage, and that the honest way to present a vug detector is to show that stage rather than the green outline in the demo.

Limitations

The circularity and area distributions, the paired estimated-vs-manual percentages, and the constant set are all drawn from the validation wells of a single engagement in one carbonate play; they are consistent across the wells and tools we tested, but a different diagenetic regime could shift the semi-circular peak or the area range, and the constants would have to be re-checked, not assumed, on new rock. The estimated-vs-manual comparison is against a single expert's per-interval picks, which are themselves an imperfect ground truth: the intervals where the pipeline caught pores the manual pass missed are evidence of that, and they cut both ways, since a manual pick that is sometimes too conservative is also sometimes too generous, and we cannot fully separate the two without a second independent interpreter. The per-interval ridge shapes in the circularity instrument are illustrative renderings of the sourced distribution rather than the exact per-interval histograms, and the redundant-merged fraction on the stage-flow lever is an illustrative reading of the duplicate-merge behaviour; both are labelled as such in the exhibits, and the distribution bounds, peak band, block area, expansion, IoU, and paired percentages they sit on are sourced. Finally, the pipeline reports an azimuth spectrum but does not interpret it; reading directional dissolution against a fracture set or a fabric is left to the geoscientist, and the tool should not be mistaken for a structural interpretation of the vug population it counts.

References

Otsu, N., 1979. A Threshold Selection Method from Gray-Level Histograms. IEEE Transactions on Systems, Man, and Cybernetics. The global-thresholding baseline whose failure on a drifting carbonate matrix motivates the adaptive, mode-subtracted approach.

<https://doi.org/10.1109/TSMC.1979.4310076>

Bradley, D. and Roth, G., 2007. Adaptive Thresholding Using the Integral Image. Journal of Graphics Tools. The local-thresholding formulation behind the block-31 Gaussian-modulated stage. <https://doi.org/10.1080/2151237X.2007.10129236>

Suzuki, S. and Abe, K., 1985. Topological Structural Analysis of Digitized Binary Images by Border Following. Computer Vision, Graphics, and Image Processing. The contour-extraction procedure that turns each thresholded copy into a candidate list.

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

Anatomy of a Top-K Mode-Subtraction Vug Detector: Five Enhanced Copies of One Log

Published June 2025. Generated 2026-07-22.

<https://www.earthscan.io/whitepapers/anatomy-of-a-top-k-mode-subtraction-vug-detector-five-enhanced-copies-one-log>

Copyright EarthScan. All rights reserved.