

One Classical-CV Parameter Set, Every Well: A Field Guide to the Eight Knobs of the Vug Detector

A classical computer-vision vug detector for carbonate borehole-image logs is attractive precisely because every decision is a named, tunable threshold rather than a weight in a checkpoint. That legibility only pays off if the parameters are portable: an operator retraining on a new well needs to know exactly which numbers travel unchanged and which must be re-established, and on what evidence. This whitepaper states the contract in full. Eight interpretable knobs govern the pipeline: top-k intensity modes ($k=5$) with a minimum mode separation ($\Delta m=5$), a centroid-difference gate ($\tau=5$ pixels), an intersection-over-union merge threshold (0.2), a Laplacian bounding-box expansion (10% each side), a mean-based enlarged radius ($2R$), a block size (31) and an adaptive-threshold constant ($C=\text{mean patch intensity}$). Six of these are universal, fixed once for every well, orientation and imaging tool; block size and C are recommended starting points; and exactly two, the Laplacian smoothness threshold and the mean-based darkness threshold, need a one-time per-well base value calibrated from one or two representative sections and then frozen for the whole well despite grayscale drift with depth. We also document two corrections that peer review forced and that materially improved the detector: a sign check on the mean difference that stops opposite-spectrum minerals being labelled vugs, and a change from deleting one of two overlapping contours to merging them, which preserves the true vug extent. The result is a rare thing in applied geoscience CV: a fully specified, reproducible pipeline whose portability is written down rather than assumed.

Quamer Nasim

ML Research Engineer

Narendra Patwardhan

Research Collaborator

VUG-DETECTION / CLASSICAL-CV / ADAPTIVE-THRESHOLDING / CARBONATE-POROSITY / BOREHOLE-IMAGE-LOGS / PARAMETER-PORTABILITY

CONTENTS

01	The portability problem, stated plainly
02	Why this is the right amount of per-well freedom
03	The mean-based filter had a sign bug, and the reviewer caught it
04	Overlapping contours: merge, do not delete
05	What the corrections say about the review
06	How to onboard a new well with this contract
07	Where a fully specified pipeline earns its keep
08	Limitations
09	References

A vug is a hole in the rock, and a vug-porosity log is a curve that a reservoir model believes. That second sentence is why the vug detector we built for a mid-sized Middle East carbonate operator is a classical computer-vision pipeline and not a segmentation network. When the porosity curve looks wrong, the petrophysicist has to be able to ask *which blob was counted, on what criterion, and why* - and get a sentence back, not a shrug. We wrote the argument for that interpretability elsewhere, in the companion piece that walks each of the eight parameters and what it physically controls, and we will not re-derive those values here; if you want block size 31 explained from the ground up, that is the piece to read.

This whitepaper answers a different and more operational question. An interpretable pipeline is only useful across a field if its parameters are *portable*, and portability is a claim that has to be earned number by number. When your team picks up the detector and points it at a well it has never seen, exactly which of the eight knobs do they leave alone, which do they treat as a starting point, and which must they re-establish before they trust a single vug? Get that wrong and the "interpretable" pipeline becomes a per-well tuning project with no end. Get it right and one parameter contract runs the whole field.

So this is the contract, stated in full, from the engagement archive. It has three parts. First, the split: six knobs are universal, two need a one-time per-well base value, and two more are recommended starting points that occasionally move. Second, a correction that peer review forced on the mean-based filter - a sign check that stops opposite-mud minerals being counted as vugs. Third, the change to overlapping-contour handling that review also forced: merge, do not delete. The first part is the spine. The other two are the honest record of what a careful reviewer caught and what we changed because they were right.

The portability problem, stated plainly

Start with what "one parameter set, every well" has to survive. Our validation spanned wells of different orientation logged on different microresistivity imaging tools - a vertical well and a second vertical well on high-resolution borehole-image logs, and a horizontal well on a compact microresistivity tool. Those wells do not present the same imagery. Grayscale levels drift with depth inside a single well as the tool and the formation change. A horizontal well images a different geometry of borehole wall than a vertical one. And two different tools have two different response characteristics that no amount of normalisation fully erases.

A naive reading says: different imagery, different parameters, tune per well. That reading is wrong, and the reason it is wrong is the whole point of the contract. Most of the pipeline's parameters are *not* image-statistics parameters at all. The number of intensity modes you subtract, the minimum spacing between them, the pixel radius at which two contour centroids are declared the same vug, the intersection-over-union at which two overlapping contours are merged, the fraction by which you expand a bounding box before measuring sharpness, the radius over which you check darkness - none of these is a property of how bright this particular well happens to be. They are properties of the *method*: how mode subtraction enhances variation, how deduplication should behave, how a false-positive filter reads a blob's neighbourhood. A method property does not care which tool acquired the log.

Only the two knobs that convert a filter's abstract sensitivity into an absolute cut on *this* well's intensity distribution are genuinely per-well. That is the entire claim, and it is small enough to write on an index card. The rest of this section makes it precise.

The six that never move

Six parameters are universal across wells, orientations and imaging tools. We fixed each one once, from analysis rather than from a per-well sweep, and never touched it again across the study:

- **k = 5**, the number of top intensity modes subtracted to build the enhanced views the detector thresholds. Five is where the marginal mode stops paying for itself; the first carries a frequency near 4,700 and the fifth near 1,700, after which subtraction enhances noise rather than features.
- **delta-m = 5**, the minimum intensity-level separation between the modes we keep. Consecutive modes (delta-m of 1) come from the same feature and produce roughly 85 to 95% duplicate contours; spacing them at five yields modes that trace materially different blobs.
- **centroid gate tau = 5 pixels**, the radius within which two contour centres are declared the same vug during deduplication.
- **IoU merge threshold = 0.2 (20%)**, the overlap at which two contours are treated as one vug.
- **Laplacian bounding-box expansion = 10%** on each side, so the sharpness filter reads the blob's rim and not just its interior.

- **mean-based radius = $2R$** , the enlarged circle over which the darkness filter compares a candidate against its surroundings.

Each of these is a decision about how the algorithm *behaves*, and behaviour does not change when the imagery does. That is why they travel.

The two that need a base value

Two parameters are genuinely per-well: the **Laplacian threshold** and the **mean-based threshold**. Both are filters that reject false positives, and both make an absolute cut on a well's intensity statistics - the Laplacian on local sharpness, the mean-based on local contrast against the matrix ring. Because grayscale distributions differ between wells and tools, the *sensitivity* of these filters is a method property but the *cut point* is not. So for each new well we establish a base value from one or two representative sections, and then - this is the operationally important half of the claim - we freeze it. It does not get re-tuned zone by zone as grayscale drifts down the well. One base value, established once, held for every depth section of that well.

The two that are starting points

Block size (31) and the adaptive-threshold constant C (the local patch mean) sit in between. They are recommended starting points rather than fixed universals, because a strongly heterogeneous well can benefit from moving them. In practice one validation well used a block size of 51 rather than 31, and one used C near -1 rather than the patch mean. But these are the exceptions that prove the rule: they moved rarely, they moved by a documented amount, and the operator always knew exactly what physical quantity they were changing - the area over which "local" is defined, or how far below the neighbourhood mean a pixel must sit to be a candidate.

One parameter set runs on every well; only two thresholds are calibrated once per well, then frozen with depth.

knobs fixed for every well

2

need a one-time base value

THE EIGHT KNOBS (PLUS THE CENTROID GATE)

top intensity modes base 5 modes	k = 5	U
mode separation base 5 levels	$\Delta m = 5$	U
centroid gate base 5 px	$\tau = 5$	U
merge threshold base 20%	IoU = 0.2	U
Laplacian window base 10% each side	bbox +10%	U
mean-based radius base enlarged 2R	r = 2R	U
Laplacian threshold base 0.25 (from sweep)	Lap. thr.	PW
mean-based threshold base 0.8 (from sweep)	mean thr.	PW

U = universal (one value everywhere) · PW = per-well base value

SWITCH WELL · WATCH ONLY THE TWO ORANGE ROWS MOVE

Horizontal well
compact imaging tool

Vertical well
borehole-image tool

Six universal knobs stay locked:

k · Δm · τ · IoU · bbox · radius

THE TWO CALIBRATED BASE VALUES

Laplacian threshold smoothness cut	0.7
mean-based threshold darkness contrast	1.0

This well also C base value re-set to -1.0 (a recommended starting point, not a re-tune of the contract).

Calibrated once from 1-2 sections, then frozen.

sourced: k=5, $\Delta m=5$, $\tau=5$, IoU 0.2, bbox +10%, r=2R, block 31, C=mean · per-well base values horizontal 0.7/1.0, vertical 0.15/0.4 · no illustrative inputs

ES-1765 · VUG PARAMETER CONTRACT UNIVERSAL VS PER-WELL · EIGHT KNOBS, TWO CALIBRATED

The classical-CV vug detector runs on one parameter contract for every well. Six knobs are universal - top modes $k=5$, mode separation $\Delta m=5$, centroid gate $\tau=5$ px, IoU merge threshold 0.2, Laplacian bounding-box expansion 10% each side, and the mean-based radius of 2R - one value each, fixed across wells, orientations and imaging tools. Exactly two knobs need a one-time per-well base value: the Laplacian smoothness threshold and the mean-based darkness threshold, each calibrated from one or two representative sections and then frozen for the whole well despite grayscale drift with depth. Toggle the horizontal compact-tool well (Laplacian 0.7, mean 1.0) against the vertical borehole-image well (Laplacian 0.15, mean 0.4) and only the two orange rows move; the six teal rows never do. Block size 31 and C=mean are recommended starting points, which is why one well re-set block to 51 and another re-set C to -1.0. Every number is sourced from the vug-paper reviewer rebuttals; there are no illustrative inputs.

The contract above is the argument of this whitepaper in one exhibit. The six teal rows are the universal knobs; the two orange rows are the per-well base values. Switch between the horizontal and vertical wells and watch what happens: the six teal rows do not move, and only the two orange thresholds change - Laplacian 0.7 and mean 1.0 on the horizontal compact-tool well, Laplacian 0.15 and mean 0.4 on the vertical borehole-image well. That is what "one parameter set, every well" actually looks like when you draw it. It is not that nothing changes between wells; it is that the surface area of what changes is two numbers, calibrated once, and everything else is a constant you can read off a card.

Why this is the right amount of per-well freedom

There is a failure mode on both sides of this contract, and the interesting engineering was finding the middle.

Fix *everything*, including the two threshold cut points, and the detector breaks the moment grayscale statistics shift enough - a filter tuned to one well's contrast will over-reject on a darker well and under-reject on a brighter one. You would be pretending that an absolute intensity cut is a method property when it plainly is not. On these logs a single pixel of the binary wireline file maps to about 3 cm of borehole, so there is already a floor of roughly ± 3 cm of depth uncertainty baked into every measurement before any filter runs; adding a mismatched threshold on top of that turns a clean pipeline into a false-positive generator.

Tune *everything* per well, and you have thrown away the reason to use a legible pipeline in the first place. Sixteen well-specific numbers is not a contract; it is a tuning habit, and it is exactly the opaque, un-auditable behaviour that a classical-CV approach was supposed to avoid. Worse, most of that tuning would be re-deriving method properties that were never well-specific to begin with, wasting an interpreter's time to reach a number that a universal constant already fixed.

The contract sits between those two, and the placement is not arbitrary. A parameter is universal exactly when it encodes how the algorithm should behave, and per-well exactly when it encodes an absolute cut on this well's intensity distribution. That is a principled line, not a convenience. It is also a *testable* one, which is why we added the two extra wells - a different orientation and a different imaging tool - specifically to prove that the six universals held and only the two thresholds needed re-establishing. All published results used those fixed base values with zero further fine-tuning, and the split survived the hardest cases we could throw at it.

THE CONTRACT IN ONE SENTENCE

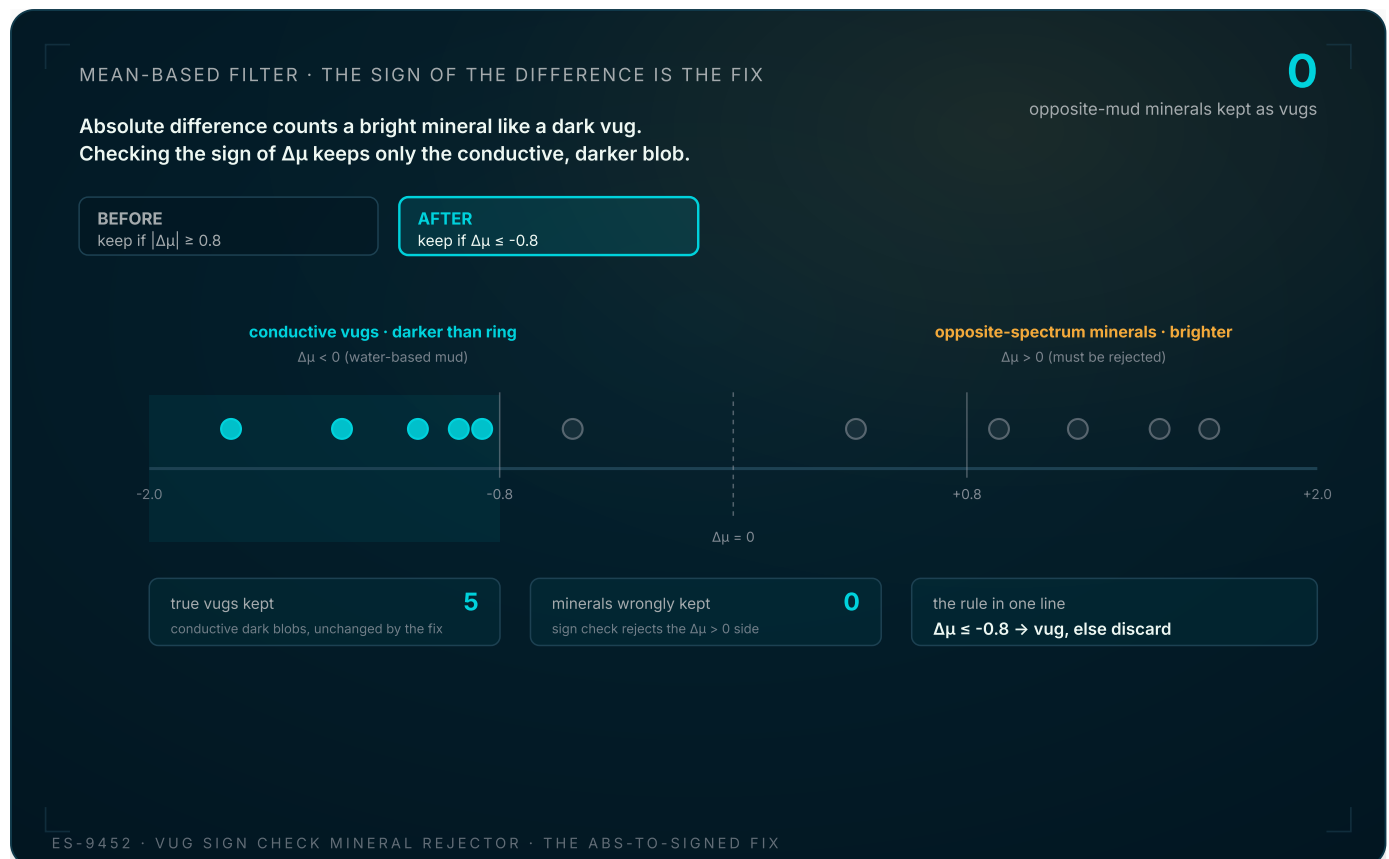
Six knobs are universal method properties and never move; block size and C are starting points that rarely do; and exactly two filter thresholds are per-well base values, each calibrated once from a representative section and then frozen for the whole well - because those two are the only numbers that encode an absolute cut on that well's own intensity distribution.

The mean-based filter had a sign bug, and the reviewer caught it

The cleanest way to earn trust in a pipeline is to show what an adversarial reader broke and what you changed. The vug paper drew a four-reviewer, roughly forty-comment rebuttal, and the single most consequential technical fix in it was a one-character conceptual error in the mean-based filter.

The mean-based filter decides whether a candidate contour is genuinely darker than the matrix ring around it. The original formulation compared the contour's mean intensity to the surrounding mean and kept the candidate when the *absolute* difference cleared a threshold. Written as an absolute value, the filter is symmetric: it keeps a blob that is much darker than its surroundings and, identically, a blob that is much brighter. On a water-based-mud log a real conductive vug is *darker* than its ring - but the opposite end of the intensity spectrum, where a resistive, bright mineral sits, has just as large an absolute difference. The absolute-value filter was letting those bright minerals through as vugs.

The reviewer flagged it, and the fix is a sign check. Define the mean difference as the surrounding mean minus the contour mean. For water-based mud, a real vug is conductive and darker, so the difference is negative; keep the candidate when the difference is at or below the threshold magnitude, and discard it when the difference is positive, because a positive difference is the opposite-spectrum mineral you never wanted. The symmetric logic holds for oil-based mud with the sign reversed. It is a small change with an outsized effect on the false-positive rate, and it is exactly the kind of error that only surfaces when a domain expert reads the equation against the physics rather than the code against the tests.



The mean-based false-positive filter keeps a contour when its mean intensity differs enough from the surrounding ring. The reviewer caught the defect: taking the absolute difference lets a feature brighter than its surroundings clear the same gate as one that is darker, so opposite-spectrum minerals (bright in water-based mud) were labelled vugs. The fix is a sign check on the mean difference: for water-based mud a real conductive vug is darker than its ring, delta-mu below zero, and delta-mu above zero is the opposite end of the spectrum and is discarded; the symmetric rule holds for oil-based mud. Toggle BEFORE (keep if the magnitude clears 0.8) against AFTER (keep only if delta-mu is at or below -0.8): before, the bright minerals on the positive side survive; after, that whole side is rejected while the true conductive vugs on the negative side are untouched. The orange markers are the mineral population the absolute-value rule wrongly kept. The rule, the 0.8 threshold, and the two-population story are sourced from the reviewer rebuttal; the individual delta-mu positions on the axis are illustrative of those two documented populations.

The exhibit above makes the bug and the fix tactile. Each blob sits at its measured mean difference along the axis; the true conductive vugs cluster on the negative side and the opposite-spectrum minerals on the positive side. Toggle BEFORE - keep if the magnitude clears 0.8 - and the bright minerals on the positive side survive as counted vugs, inflating the porosity. Toggle AFTER - keep only if the difference is at or below -0.8 - and that whole side is rejected while the genuine vugs on the negative side are untouched. The orange markers are precisely the mineral population the absolute-value rule wrongly kept. Nothing about the true vugs changed; the fix cost nothing on the signal and removed a class of false positive wholesale.

It is worth being clear about why this matters beyond one paper. A porosity number that includes mislabelled minerals is not slightly wrong - it is wrong in a direction that biases the reservoir model toward more secondary porosity than the rock has. An absolute-value filter would have shipped that bias silently. A legible filter, read by a reviewer who knew that conductive and resistive features sit at opposite ends of the spectrum, did not.

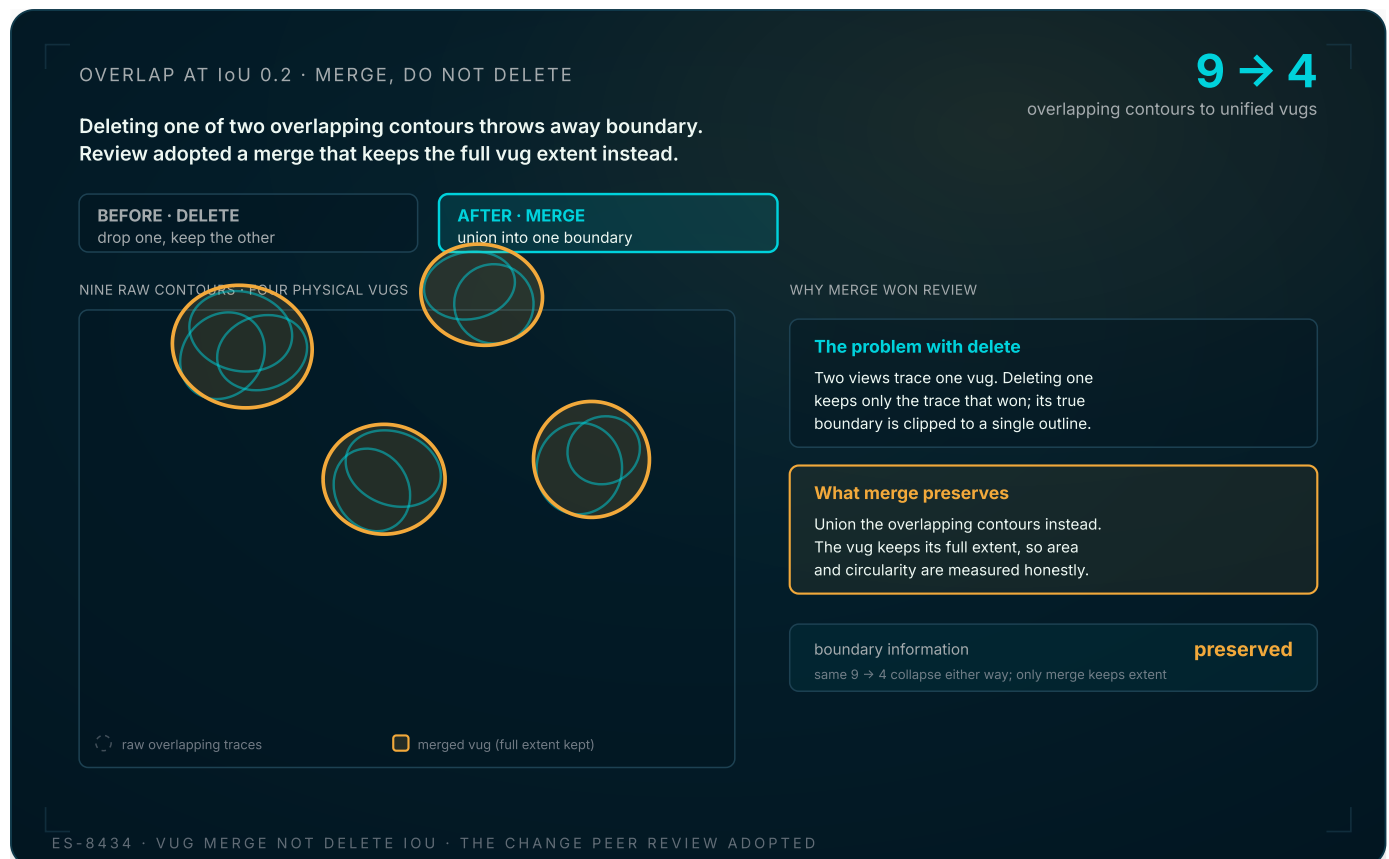
Overlapping contours: merge, do not delete

The second correction review forced concerns deduplication, and it is a good illustration of how a defensible default can still be the wrong one.

Because the detector subtracts five intensity modes and thresholds each, the same physical vug is often traced up to five times, once per enhanced view. Those redundant contours have to be collapsed or the porosity is wildly inflated. The pipeline collapses them with two readable numbers - the centroid gate ($\tau = 5$ pixels) and the IoU merge threshold (0.2) - that together decide when two contours describe one vug. The question review sharpened was what to *do* once you have decided two contours are the same vug.

The original pipeline deleted one of each overlapping pair and kept the other. That is defensible: you end with one contour per vug, the count is correct, and it is the obvious analogue of the non-maximum suppression a learned detector runs. But it quietly discards information. The deleted contour was a real trace of the same blob from a different enhanced view, and its boundary carried part of the vug's true extent. Keeping only the surviving contour clips the vug to whichever single trace won, and because area and circularity are computed from that boundary, a clipped boundary means a mis-measured vug.

The reviewer's suggestion, which we adopted, was to *merge* the overlapping contours into one unified boundary rather than delete either. The documented demonstration collapsed nine overlapping contour regions to four unified vug boundaries, with each merged outline preserving the full extent that a delete would have clipped. Same count - nine contours become four vugs either way - but the merged boundary measures the vug honestly.



Five mode-subtracted views trace the same physical vug several times, so contours whose intersection-over-union clears 0.2 must be collapsed to one. The original pipeline deleted one of each overlapping pair, which discarded the boundary the deleted contour carried. The change peer review actually adopted was to merge the overlapping contours into one unified boundary instead of deleting either, so the vug keeps its full extent and its area and circularity are measured honestly. The documented demo collapsed nine overlapping contour regions to four unified vug boundaries. Toggle BEFORE (delete: keep one contour per cluster, drop the rest) against AFTER (merge: union each cluster): both collapse nine to four, but only merge preserves the true extent, shown as the orange boundary the delete rule would clip away. The merge-not-delete change, the nine-to-four count, and the 0.2 threshold are sourced from the reviewer rebuttal; the specific contour polygons are an illustrative reconstruction of that demo.

The exhibit above runs both rules on the same nine contours. Under BEFORE·DELETE, each cluster keeps one contour and drops the rest, and the surviving outline is whichever single trace won. Under AFTER·MERGE, each cluster is unioned into one boundary. Both collapse nine to four; the orange outline is the extent that merge preserves and delete throws away. The verdict readout states the difference in the terms that matter downstream: the count is identical, but only merge keeps the boundary information that area and circularity are measured from.

The general lesson is one worth carrying into any deduplication step, learned or classical. Deleting a duplicate is not free when the duplicate is a partial view of the same object; the union of the partial views is a better estimate of the whole than any single survivor. Non-

maximum suppression's instinct - keep the best, kill the rest - is the wrong instinct when "the rest" are complementary traces rather than redundant guesses.

What the corrections say about the review

Both fixes came from reviewers reading equations against physics, and both improved the detector rather than merely satisfying a comment. That is the argument for publishing a fully specified classical pipeline in the first place: a legible method can be audited by a domain expert, and a domain expert audit catches errors that a purely computational test suite never would. A segmentation network that mislabelled resistive minerals as vugs would have failed the same way, but there would have been no equation to read and no sign to check - only a distribution of weights and a worse number. The interpretability is not a nicety layered on top of the method. It is the property that let the method be corrected.

There is a related discipline the review reinforced, worth naming because it is easy to skip. Missing pixels on these logs are coded as a sentinel value and must be converted to a genuine null and excluded from every filter - the Gaussian modulation, the Laplacian, and the mean-based comparison - or they poison the local statistics that all three depend on. A visualisation bug once rendered those nulls as if they were data in one figure; the computation was always correct, but the figure was not, and the fix was to mask the nulls in the display too. It is a small thing that stands for a large one: in a pipeline whose output a reservoir model trusts, the gap between "the computation is right" and "what we showed is right" is itself a defect to close.

How to onboard a new well with this contract

Put the pieces together and the operational procedure for a well the detector has never seen is short, which is the entire point.

Leave the six universal knobs exactly as they are: $k = 5$, $\Delta m = 5$, centroid $\tau = 5$, IoU 0.2, Laplacian bounding-box expansion 10%, mean-based radius $2R$. Start block size at 31 and C at the local patch mean, and move them only if the well is heterogeneous enough to warrant it, in which case document the move and its reason - a block of 51 covers a larger physical area over which "local" is defined; a C near -1 shifts how far below the neighbourhood mean a pixel must sit. Then do the one piece of real per-well work: pick one or two representative sections, establish the Laplacian and mean-based base values there,

confirm the detector's vugs on those sections against the interpreter's reading, and freeze both thresholds for the rest of the well. Do not re-tune them as grayscale drifts with depth; the base value is a property of the well, not of the zone.

That is the whole procedure. One well, two numbers to establish, six constants to leave alone, two starting points to consider. It is short because the contract did the hard thinking once, up front, and wrote down which numbers are allowed to move.

Where a fully specified pipeline earns its keep

The reason to invest in stating a contract this precisely is that it is the difference between a result and a tool. A result is a number in a paper produced by a configuration someone remembers. A tool is a pipeline an operator's own team can run on next year's wells without the original authors in the room, and it only becomes that when the portability is written down: which parameters are constants of the method, which are properties of the well, and how you establish the latter without a sweep.

For a petrophysicist or an R&D geoscience lead evaluating this class of detector, the questions that separate the two are concrete:

- For every parameter, is it stated whether it is universal, a starting point, or per-well - and is the per-well set small enough to establish without an open-ended tuning project?
- Are the per-well base values established from a defined, minimal set of representative sections and then frozen, or re-tuned zone by zone?
- Does each false-positive filter encode a physical claim you can read and argue with, and has the sign of that claim been checked against the mud type?
- When two contours describe the same vug, does the pipeline preserve the vug's extent, or clip it to a single survivor?
- Is every reported number read against the instrument's physical resolution floor rather than in spite of it?

This detector was built to answer yes to all five, and the reason it can is that the interpretability was never decorative. Eight knobs, six of them fixed for the whole field, two calibrated once per well, two starting points that rarely move, and two corrections a careful review made better. That is not a compromise against a deep network. On a QC deliverable

that a domain expert must trust, sign off on, and feed into a model that decides where to drill, a pipeline whose portability is a written contract is the stronger tool - because it can be read, audited, corrected, and run by someone who was not there when it was built.

WHAT THIS WHITEPAPER ARGUES

- 1 A legible classical-CV vug detector is only useful across a field if its parameters are portable, and portability is a claim that must be earned parameter by parameter, not assumed.
- 2 Six of the eight knobs are universal method properties fixed for every well, orientation and imaging tool ($k=5$, $\Delta m=5$, centroid $\tau=5$ px, IoU 0.2, Laplacian bbox +10%, mean-based radius $2R$); block size 31 and $C=\text{mean}$ are recommended starting points that rarely move.
- 3 Exactly two knobs are genuinely per-well: the Laplacian and mean-based thresholds, each an absolute cut on the well's intensity distribution, calibrated once from one or two representative sections and then frozen for the whole well despite grayscale drift with depth (horizontal well 0.7/1.0, vertical well 0.15/0.4).
- 4 A reviewer caught a sign bug: the mean-based filter's absolute difference let opposite-spectrum minerals (bright in water-based mud) pass as vugs; the fix is a sign check on the mean difference (keep only the conductive, darker side), with symmetric logic for oil-based mud.
- 5 Overlapping-contour handling was changed from delete to merge during review: deleting one of two overlapping contours clips the vug's extent, while merging nine overlapping regions into four unified boundaries preserves the extent that area and circularity are measured from.

Limitations

This contract is validated on three wells from one carbonate play, spanning two orientations and two imaging tools; three wells is enough to demonstrate that the universal-versus-per-well split holds across those axes, but it is not enough to bound how the two per-well thresholds vary across a full field, and an operator onboarding many wells should treat the base-value calibration as a step to monitor rather than a solved problem. The two threshold

values reported for the horizontal and vertical wells are the base values those wells used, not a range; we make no claim that they generalise numerically to other wells, only that the *procedure* for establishing them does. The block-size and C exceptions are recommendations rather than a rule, and a sufficiently heterogeneous well may move parameters we have treated as universal - the contract states the common case, and documents its own exceptions, but it is not a proof. The depth resolution of the input imposes a hard floor of roughly ± 3 cm on any depth-referenced measurement, independent of every parameter here. And the whole pipeline is unsupervised classical CV: it carries no learned prior about what a vug looks like beyond the eight named criteria, which is the source of both its auditability and its ceiling. The individual contour polygons and mean-difference marker positions shown in the exhibits are illustrative reconstructions of the documented behaviour, not per-blob measurements; the parameter values, the sign-check rule, the merge-not-delete change, and the nine-to-four collapse are the sourced facts.

References

[1] Suzuki, S., and Abe, K. (1985). Topological Structural Analysis of Digitized Binary Images by Border Following. Computer Vision, Graphics, and Image Processing. The border-following algorithm the pipeline uses to extract candidate vug contours.

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

[2] Bradley, D., and Roth, G. (2007). Adaptive Thresholding Using the Integral Image. Journal of Graphics Tools. The local-mean adaptive-threshold formulation underneath the block-size and C constant. <https://doi.org/10.1080/2151237X.2007.10129236>

[3] Otsu, N. (1979). A Threshold Selection Method from Gray-Level Histograms. IEEE Transactions on Systems, Man, and Cybernetics. The global-threshold baseline that per-neighbourhood adaptive thresholding replaces on zoned carbonate imagery.

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

One Classical-CV Parameter Set, Every Well: A Field Guide to the Eight Knobs of the Vug Detector

Published May 2025. Generated 2026-07-22.

<https://www.earthscan.io/whitepapers/one-classical-cv-parameter-set-every-well-eight-knobs-vug-detector>

Copyright EarthScan. All rights reserved.