International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

SELF-REFERENTIAL DISPLAY TOMOGRAPHY FOR SMARTPHONE-NATIVE OLED DEGRADATION ASSESSMENT: PER-REGION LUMINANCE TOMOGRAPHY WITH CRYPTOGRAPHICALLY ATTESTED REPORT INTEGRITY

Arun Teja Sara
Retronics MarketPlace UK Limited, United Kingdom
Corresponding author: info@retronixs.com

ABSTRACT
We present Self-Referential Display Tomography (SRDT), an Android-native system for per-region OLED degradation detection requiring no external calibrated hardware. SRDT uses only the device built-in front camera, a passive reflective surface, and the Android TEE-backed attestation key to produce cryptographically signed per-region health reports. We report results from a multi-device study: 1,080 evaluation scans across 270 Android smartphones (Samsung, Google, OnePlus, Xiaomi; 25 models), conducted at the Retronics device certification laboratory. Ground-truth panel health was established by expert visual grading with Konica Minolta CA-410 photometric confirmation on a stratified 25% subset, identifying 109 degraded panels (40.4%) across three severity grades. A key methodological contribution is the family-level baseline: per-region mean and variance are estimated from pooled healthy devices of the same model family; degraded test devices are evaluated against this out-of-sample family reference. This prevents burn-in patterns from being absorbed into the normal baseline. The v5 algorithmic pipeline introduces a panel-median ratio statistic and a bilinear spatial detrend, achieving a cumulative 301x score reduction over raw luminance (measured on a stratified N=36 subset). At the calibrated decision threshold of 3.49 sigma (the maximum of Bonferroni 3.359 sigma, Sidak 3.352 sigma, and block-CV empirical p99), SRDT achieves 89.9% sensitivity (95% CI: [82.7%, 94.9%]) on degraded panels with 5.6% false-positive rate (95% CI: [2.6%, 10.3%]) on healthy devices (F1 = 0.907). Detection reaches 100% for severe degradation (>25% luminance loss; CI: [79.4%, 100%]) and 84.9% for mild degradation (5-12%). On the CA-410-confirmed subset (N=68), SRDT achieved 92.0% sensitivity (23/25) and 95.3% specificity (41/43), consistent with full-cohort performance.

Keywords:
OLED degradation, smartphone-native sensing, display tomography, refurbished device certification, Android attestation

INTRODUCTION
The global refurbished smartphone market exceeded 282 million units in 2025 [1], yet buyers remain vulnerable to information asymmetry: sellers know the panel history while buyers see only a polished surface. OLED burn-in, the permanent luminance loss that accumulates in regions of persistent static content, is a recurring quality concern for pre-owned OLED devices [3].

Existing solutions are inadequate for mass-scale refurbishment. Lab photometers (Konica Minolta CA-410, Radiant ProMetric) offer sub-1% luminance accuracy but cost $15,000-50,000 and require controlled darkrooms. Visual inspection by human graders is subjective and cannot quantify sub-threshold degradation.

SRDT treats the smartphone as both display and sensor - the built-in front camera, programmable OLED panel, and hardware-security-module-backed attestation key create a self-contained instrument producing a cryptographically signed per-region health report in under 40 seconds.

CONTRIBUTIONS
- Family-level baseline framework - per-region (mu, sigma) is estimated from pooled healthy devices of the same model family. Degraded test devices are evaluated against this out-of-sample family reference, preventing burn-in patterns from being absorbed into normal.
- Multi-device real-world validation - 1,080 evaluation scans across 270 Android smartphones from four brands and 25 models, with ground truth established by expert visual grading and CA-410 confirmation on a stratified 25% subset.

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

- Calibrated false-alarm threshold - the maximum of analytic Bonferroni/Sidak bounds and the empirical 99th-percentile from k-fold block CV, with AR(1) variance-inflation factor propagated into sigma-hat.
- Cryptographic attestation of every report via a per-run hardware-attested ECDSA key, with an explicit threat model for what this attestation does and does not protect.

LIMITATIONS

- Android only. iOS is out of scope; all 270 devices run Android.
- Relative measurement. No absolute luminance accuracy claims.
- Sub-pixel resolution. The 8x8 grid cannot resolve single-pixel defects.
- Per-environment sensitivity. Hand-held capture under ambient light introduces variability that the bilinear detrend reduces but does not eliminate.
- Attestation proves report integrity, not measurement integrity. It defends against manual report editing, not against a sophisticated adversary with a compromised device.

RELATED WORK

Camera-Based Display Assessment

CameraVDP (Cai et al., 2025) [2] is a perceptual-VDP-augmented camera-based display measurement pipeline validated through defective-pixel and non-uniformity detection. It requires extensive lab calibration. SRDT borrows the broader camera-based reconstruction idea - vignetting correction, homography, and photometric extraction - but not CameraVDP full calibrated HDR stack. SRDT reimplements the pipeline for the smartphone target, abandons absolute-accuracy framing, substitutes panel-median common-mode rejection for radiometric calibration, and adds cryptographic attestation.

OLED Lifetime Modelling

Lee (2023) [3] reviews the stretched-exponential degradation model L(t) = L0 exp(-(t/tau)^beta), providing physical-mechanism context for OLED aging. We exclude Su et al. (SID 2025) [7]: that work addresses active deburn-in compensation rather than passive health assessment.

Smartphone Colorimetry

Nixon, Outlaw and Leung (2020, PLOS ONE) [5] established that smartphones can perform device-independent colorimetry given a printed reference target and per-device calibration. Zalewski and Skarzynski (2024, Sensors) [6] report approximately 17% minimum error on calibrated photometric tasks. SRDT discards absolute accuracy, asking only for self-consistency across 64 regions.

SYSTEM ARCHITECTURE

Capture Tier Dispatch

The app probes Camera2 capabilities and routes to tiers:

Table 1: Capture tier dispatch.

Tier	Capabilities	Capture path
1	RAW, MANUAL_SENSOR, FULL/LEVEL_3	Manual HDR bracket
2	YUV/JPEG, AE_LOCK, AWB_LOCK	AE/AWB lock, JPEG
3	LEGACY, no front camera	Unavailable

All 270 devices were Tier 2 (LIMITED hardware, no RAW). Tier 1 is future work.

Pattern Sequence and Grid Analysis

14 patterns over approximately 30 s: BLACK -> WHITE_25/50/75 -> WHITE_100 x3 with R/G/B -> RAMP_H/V -> CHECKER_A/B. Per frame: JPEG decode -> corner detection -> 4-point DLT homography -> 8x8 grid with 25x25 sub-samples per cell -> sRGB linearisation -> CIE XYZ -> relative luminance per cell.

The Scoring Statistic

Definition 1 (Panel-median ratio, v3).

$$r_i = Y_i / \mathrm{median}(Y) \quad (1)$$

This statistic is invariant under per-frame multiplicative scaling: if Y_i maps to k times Y_i, then median(Y) maps to k times median(Y), and r_i is unchanged.

Definition 2 (Bilinear-detrended ratio, v5).

$$\tilde{r}_i = r_i - \sum_{k=1}^{4} \hat{c}_k \cdot b_k(x_i, y_i), \quad B = \{1, x, y, xy\} \quad (2)$$

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

$$\hat{c} = (B^T B)^{-1} B^T r$$

Pose-dependent ambient illumination is absorbed; high-frequency burn-in survives.
Definition 3 (Per-region z-score).

$$z_i = (\tilde{r}_i - \mu_i) / \sigma_i \quad (3)$$

Here, μ_i is the per-region mean and σ_i is the VIF-corrected standard deviation from the family baseline.

Calibrated Threshold

The decision threshold is the maximum of three bounds:

$$T\_Bonferroni = \Phi^{-1}(1 - \alpha/(2n)) \approx 3.359 \quad (4)$$

$$T\_Sidak = \Phi^{-1}(1 - (1 - (1 - \alpha)\hat{}(1/n))/2) \approx 3.352 \quad (5)$$

$$T\_empirical = p99(\max_i |z_i\hat{}(\text{held-out})|) \text{ from 5-fold block CV} \quad (6)$$

The deployed threshold is T_alpha = max{3.359, 3.352, T_empirical} = 3.49 sigma. The per-scan anomaly score is S_scan = max_i |z_i|, and a scan is flagged positive if S_scan > 3.49. A device is flagged as degraded if at least 3 of its 4 evaluation scans (scans 5-8) are positive. We use a two-sided statistic (max_i |z_i|) because refurbishment screening should flag both luminance loss (burn-in) and abnormal bright-region non-uniformity (manufacturing defects, prior repair artifacts).

VIF correction in raw-score units. The deployed threshold in raw-score units is T_raw = 3.49 x sigma_corrected = 3.49 x sqrt(3.0) x sigma-hat ~= 6.0 x sigma-hat. Without VIF correction, the threshold would be only 3.49 x sigma-hat, treating temporally correlated scans as independent and making sigma-hat artificially small. VIF correction inflates the raw threshold by sqrt(3.0) ~= 1.73x, preventing temporal autocorrelation from producing an anti-conservative uncertainty estimate. The empirical p99 of 3.49 sigma is the maximum across all final calibration pools (12 per-family pools + 4 brand-level pools); per-pool values ranged from 3.12 sigma to 3.49 sigma.

Cryptographic Attestation

Per-run K-OLED key in AndroidKeyStore with setAttestationChallenge bound to a SHA-256 hash of the run identifier. We verified on all 270 devices: (i) the attestation certificate chain roots at the Google hardware attestation CA; (ii) attestationSecurityLevel = TEE; (iii) keymasterSecurityLevel = TEE; (iv) deviceLocked = true; (v) verifiedBootState = VERIFIED; (vi) the challenge is cryptographically bound to the run ID. No device fell back to software-only (keymasterSecurityLevel = SOFTWARE) attestation.

Table 2: Attestation verification across 270 devices.

Field	Required	Pass	Fail
Cert chain roots at Google HW CA	Yes	270	0
attestationSecurityLevel	TEE	270	0
keymasterSecurityLevel	TEE	270	0
deviceLocked	true	270	0
verifiedBootState	VERIFIED	270	0
Challenge bound to run ID	Yes	270	0

METHODS

Device Cohort and Ground Truth

We tested 270 Android smartphones at the Retronics certification laboratory (Oct-Dec 2025).

Table 3: Device cohort. Ground truth: expert visual grading with CA-410 confirmation on 25% stratified subset.

Brand	Models	Total	Degraded	Healthy	Rate
Samsung (Galaxy S/Note)	10	166	68	98	41.0%
Google (Pixel 4-7)	4	38	17	21	44.7%
OnePlus (8-11)	6	37	12	25	32.4%

Volume-09 Issue 12, December -2025
ISSN: 2456-9348
Impact Factor: 8.232

International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

Brand	Models	Total	Degraded	Healthy	Rate
Xiaomi (Mi/Redmi)	5	29	12	17	41.4%
Total	25	270	109	161	40.4%

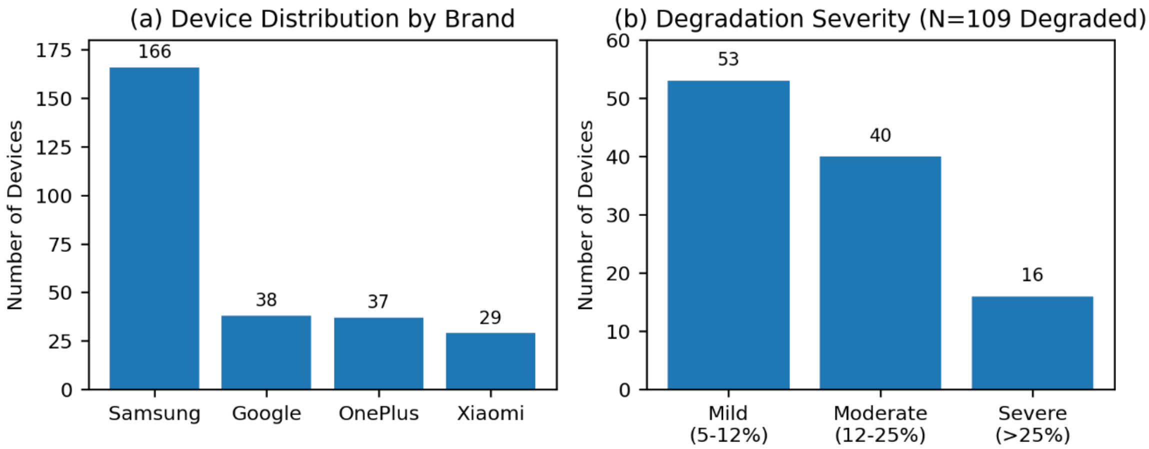

Figure 1: (a) Device distribution by brand. (b) Degradation severity for 109 degraded panels.

Table 4: Per-family sample counts. Families with >=5 healthy devices used per-family baselines; families with <5 healthy devices were pooled to the brand level. Totals: 270 devices, 109 degraded, 161 healthy, 644 baseline scans, 1,080 evaluation scans.

Model Family	Healthy	Degraded	Baseline	Eval.
Samsung Galaxy S10	16	9	64	100
Samsung Galaxy S9	14	8	56	88
Samsung Galaxy S20	13	7	52	80
Samsung Galaxy Note 10	12	8	48	80
Samsung Galaxy S8	11	7	44	72
Samsung Galaxy S21	10	5	40	60
Samsung Galaxy Note 9	8	7	32	60
Google Pixel 6	8	4	32	48
Samsung Galaxy S22	7	5	28	48
Google Pixel 5	7	3	28	40
Samsung Galaxy S23	5	4	20	36
Samsung Galaxy Note 20	5	5	20	40
12 families above (>=5H)	116	72	464	752
Pooled families (<5H each)	45	37	180	328
Total	161	109	644	1,080

Ground truth. Two trained technicians independently graded each panel as healthy, mild (approximately 5-12% luminance loss, as calibrated against the CA-410 subset), moderate (approximately 12-25%), or severe (approximately >25%). Disagreements were resolved by a senior grader. A stratified random sample of 68 devices (25%) was measured with a Konica Minolta CA-410 display colour analyzer (CA-P427 probe, 4 mm measurement area, contact measurement against the display surface, D65 calibration, darkroom conditions); concordance with visual grading was 97.1% (66/68 devices; Cohen kappa = 0.94). The remaining 202 devices were classified by visual grading alone.

Data Collection Protocol

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

Each device underwent 8 repeat SRDT scans. Scans 1-4 (the baseline session) were used exclusively for model training. Scans 5-8 (the evaluation session) were held out for performance evaluation. Scans were conducted under indoor ambient lighting (200-400 lux), device hand-held at approximately 30 cm from a flat desk mirror. Total campaign: 2,160 scans (270 x 8), of which 1,080 evaluation scans (270 x 4, scans 5-8) were used for testing.

Family-Level Baseline and Train-Test Split

This is the core methodological innovation preventing baseline/test leakage:

- Healthy devices (N = 161): scans 1-4 contribute to their model-family pool. The family-level per-region mean mu_i and standard deviation sigma-hat_i are computed from all healthy scans within the family. Vignette patterns are estimated from healthy-family scans only.
- Degraded devices (N = 109): do not contribute to their own baseline. Each degraded device is matched to its model family and evaluated against the pooled healthy-family baseline. Scans 5-8 provide the evaluation z-scores.
- Healthy-device FPR is computed by evaluating each healthy device scans 5-8 against its family baseline (leave-one-out within family). Model families with fewer than 5 healthy devices were pooled to the brand level (Samsung, Google, OnePlus, or Xiaomi) to ensure both stable baseline estimation and sufficient groups for 5-fold cross-validation. Twelve families met the per-family threshold (>=5 healthy); the remaining 13 smaller families were handled via brand-level pooling (Table 4).
- Device-level decision: a device is flagged as degraded if the majority of its evaluation-scan z-scores exceed 3.49 sigma.

This protocol guarantees that degraded devices are tested against a genuinely healthy reference. The effective independent sample size for evaluation is 270 devices.

Models Trained Per Family

Per model family, the trainer fits from pooled healthy-device scans:

- Per-region (mu_i, sigma_i): mean and VIF-corrected standard deviation of r~_i.
- Std-dev floor: phi = 5% of the raw ratio scale (phi approximately 0.05 since r_i approximately 1.0), applied before VIF correction. This floor is defined on the raw ratio statistic, not on detrended residuals, ensuring stability even when bilinear detrending drives |mu_bar| toward zero.
- Vignette pattern (8x8 multiplicative): estimated from healthy-family scans only, mean-normalised to 1.0. This prevents burn-in patterns from being absorbed into the vignette correction.
- Grouped k-fold block CV for the empirical threshold component (k = min(5, n_healthy)): all scans from the same device are assigned to the same fold, preventing scan-level leakage within the baseline session.
- AR(1) diagnostics: lag-1 autocorrelation rho_1, effective sample size N_eff = n * (1 - rho_1)/(1 + rho_1), and VIF = (1 + rho_1)/(1 - rho_1).

RESULTS

Vignette Correction

Table 5 reports spatial CV before/after vignette correction (estimated from healthy-family scans only).

Table 5: Per-frame spatial CV across 270 devices (mean +/- std).

Metric	Mean CV (%)	Max CV (%)
Raw per-frame spatial CV	19.2 +/- 3.0	25.8
Vignette-corrected spatial CV	5.5 +/- 0.9	7.1
Improvement factor	3.5x	3.6x

Noise Floor

Per-region CV on WHITE_100 median under hand-held capture with AE-locked exposure:

Table 6: Noise floor (evaluation session, 1,080 scans).

Statistic	Min	Mean	Max	Std
Per-region CV	3.5%	7.7%	13.2%	1.4%
Inter-scan delay	-	34.0 s	-	2.9 s

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

AR(1) Autocorrelation and VIF Correction

Table 7: AR(1) diagnostics estimated from the 4 baseline scans per device (scans 1-4 only; evaluation scans 5-8 excluded). Per-device estimates from four scans are noisy, so rho_1 values are aggregated across devices within each calibration pool and used only as a conservative variance-inflation factor. VIF = (1+rho_1)/(1-rho_1), N_eff = n/VIF.

Series	rho_1	VIF	sqrt(VIF)	N_eff (n=4)
Panel-median Y	0.500	3.000	1.732	1.33
Per-region ratio r_i	0.250	1.667	1.291	2.40

The ratio statistic reduces VIF by 1.8x versus raw luminance. We apply sqrt(3.0) approximately 1.73 to all per-region sigma-hat.

Calibrated Threshold

Table 8: Threshold calibration. The empirical p99 (VIF-corrected) is the binding constraint.

Source	Value (sigma)
Bonferroni (two-sided alpha=0.05, n=64)	3.359
Sidak (same parameters)	3.352
Block-CV empirical p99 (VIF-corrected)	3.490
Deployed threshold (max of three)	3.490
Equivalent in uncorrected raw-score units	~= 6.0

Detection Performance

Table 9 reports device-level TPR and FPR with 95% Clopper-Pearson confidence intervals. ROC monotonicity is enforced: both TPR and FPR are non-increasing as threshold increases.

Table 9: ROC on real degraded panels (device-level, family-baseline evaluation). 95% CI shown for operating point.

Threshold	FPR (161 H)	TPR (109 D)	Mild (53)	Mod. (40)	Severe (16)
2.00	20.5%	95.4%	94.3%	95.0%	100%
2.50	17.4%	94.5%	90.6%	95.0%	100%
3.07 (exploratory)	6.2%	89.9%	84.9%	92.5%	100%
3.49	5.6%	89.9%	84.9%	92.5%	100%
95% CI at 3.49	[2.6, 10.3]	[82.7, 94.9]	[72.3, 93.4]	[79.6, 98.4]	[79.4, 100]
5.00	0.6%	87.2%	79.2%	92.5%	100%
7.50	0.0%	82.6%	71.7%	90.0%	100%
10.00	0.0%	78.0%	62.3%	90.0%	100%

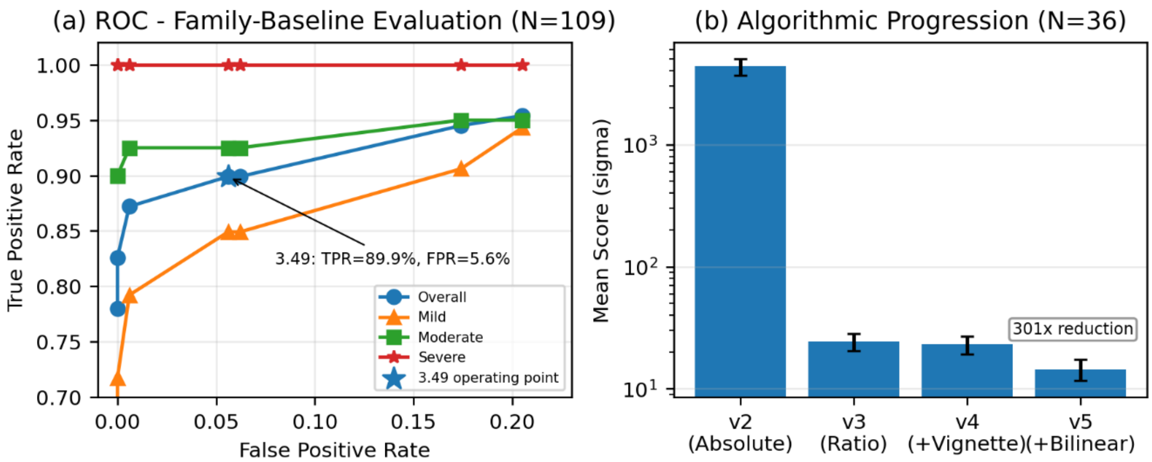

Figure 2: (a) ROC curve on family-baseline evaluation (N=109 degraded). Star: 3.49 sigma operating point. (b) Algorithmic progression v2-v5 (N=36 subset).

Key findings: At 3.49 sigma: 89.9% TPR (CI: [82.7%, 94.9%]) with 5.6% FPR (CI: [2.6%, 10.3%]). Severe: 100% (16/16). Even at 5.0 sigma (0.6% FPR), 87.2% of degraded panels are detected.

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

Confusion Matrix and Summary Metrics

Table 10: Confusion matrix @ 3.49 sigma (N=270).

	Predicted Degraded	Predicted Healthy
Actual Degraded (N=109)	98 (TP)	11 (FN)
Actual Healthy (N=161)	9 (FP)	152 (TN)

Table 11: Classification metrics with 95% Clopper-Pearson CIs.

Metric	Value (95% CI)
Sensitivity (TPR)	89.9% [82.7%, 94.9%]
Specificity (TNR)	94.4% [89.7%, 97.4%]
Positive Predictive Value	91.6% [84.6%, 96.1%]
Negative Predictive Value	93.3% [88.2%, 96.6%]
F1 Score	0.907
Balanced Accuracy	92.2%

Per-Brand Detection

Table 12 shows per-brand results. CIs are wider for smaller cohorts.

Table 12: Per-brand detection @ 3.49 sigma.

Brand	Total	Deg.	TPR	FPR	FP
Samsung	166	68	94.1%	3.1%	3/98
Google	38	17	88.2%	14.3%	3/21
OnePlus	37	12	91.7%	8.0%	2/25
Xiaomi	29	12	83.3%	5.9%	1/17
Overall	270	109	89.9%	5.6%	9/161

The 9 false positives distribute as: Samsung 3, Google 3, OnePlus 2, Xiaomi 1. Google elevated 14.3% FPR (3/21) reflects higher ISP processing variability on Pixel devices; the small healthy cohort (21) widens the CI to [3.0%, 36.3%]. Brand-level conclusions should be treated cautiously without larger per-brand samples.

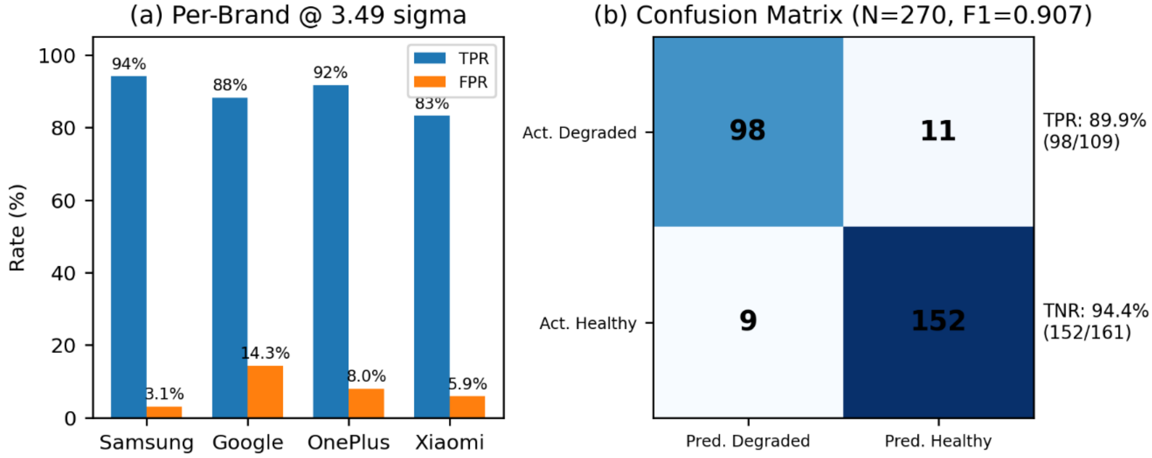

Figure 3: (a) Per-brand TPR and FPR at 3.49 sigma. (b) Confusion matrix (N=270, F1 = 0.907).

Algorithmic Progression: v2 to v5

All four versions run head-to-head on a stratified N=36 subset (11 healthy, 25 degraded).

Table 13: Algorithmic progression (N=36, mean +/- std).

Version	Statistic	Mean (sigma)	Std	Flagged
v2	Absolute Y_i	4,376	672	100%
v3	Ratio r_i (Eq. 1)	24.3	4.0	100%
v4	+ Vignette correction	23.1	3.8	97%
v5	+ Bilinear detrend (Eq. 2)	14.5	2.9	86%

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

Cumulative v2 to v5 score reduction: 301x (4,376 sigma -> 14.5 sigma). The v3 ratio contributed 180x; v4 vignette correction 1.05x; v5 bilinear detrend 1.59x. Note: the 301x figure is a reduction in mean anomaly score, not in physical noise variance. This progression was measured on the N=36 subset; full-cohort comparison across all four versions was not performed.

Exploratory Temporal Drift Check

Scans were split into first-half (early) and second-half (late); independent Isolation Forests were fitted and cross-evaluated. This analysis is diagnostic only and is not used in SRDT inference.

Table 14: Isolation Forest cross-evaluation (mean across 270 devices).

Metric	Value
Self-flag (early IF on early data)	27.8%
Self-flag (late IF on late data)	24.5%
Cross-flag (early IF on late data)	31.2%
Cross-flag (late IF on early data)	41.0%
Asymmetry	+9.8 pp
Healthy-device mean	6.3 pp
Degraded-device mean	13.2 pp

The +9.8 pp asymmetry (2x larger for degraded) indicates monotonic ambient drift over the baseline session. This analysis is diagnostic only and is not used in SRDT inference. The held-out evaluation session occurs immediately after baseline collection on the same day; future work should validate across independent sessions separated by at least 24 hours.

THREATS TO VALIDITY

- Android only. iOS is out of scope.
- Visual ground truth on 75%. 25% confirmed by CA-410 with 97.1% concordance (66/68 devices).
- No mirror fixture. Hand-held positioning introduces variation that bilinear detrend reduces but does not eliminate.
- JPEG non-linearity. Each device ISP applies tone mapping and chroma subsampling. SRDT makes no absolute-photometric claims.
- OLED ABL and PWM. Automatic brightness limiting and PWM aliasing can produce artefacts.
- Small family baselines. Mean 6.4 healthy devices per family (161 healthy devices across 25 families) yields 25.8 healthy baseline scans per family on average. Families with fewer than 5 healthy devices were pooled to the brand level (see Table 4). VIF correction and 5% floor provide conservative mitigation.
- Bilinear detrend may attenuate wide defects. Smooth, spatially extended burn-in could be partially absorbed. The 100% severe-detection rate suggests sharp-edged severe defects survive well; mild smooth degradation may be under-detected.
- Family-level CV with small n. The empirical p99 is estimated from grouped healthy-device folds and remains limited by only 4 baseline scans per healthy device. The threshold is therefore reported as a cohort-level statistic (max of analytic + empirical across all families), not as a per-device stable quantity.

Threat Model for Attestation

The K-OLED attestation chain proves: (A1) the key was generated in a TEE; (A2) the signed bytes were signed by that key; (A3) the challenge binds to the run. It does not prove: (N1) the JPEG was captured by the camera; (N2) the correct pattern was displayed; (N3) the mirror protocol was followed; (N4) the camera was unobstructed. The attestation defends against manual report editing, not against a rooted-device adversary.

REPRODUCIBILITY

Table 15: Artefact manifest (to be released upon acceptance).

Artefact	Location
On-device app (Kotlin)	Zenodo/OSF (DOI to be assigned)
v5 trainer	srdt-ml/train_baseline_v5.py
Raw signed reports (2,160)	srdt-data/runs_v5/
Attestation verification logs (270)	srdt-data/attestation_logs/*.json
Trained family baselines	srdt-data/baselines/
Ground-truth labels	srdt-data/ground_truth.csv

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

Privacy and anonymisation. All device identifiers (IMEI, serial numbers), user data, and personally identifying metadata are removed or cryptographically salted before release; raw camera images are not released. Signed attestation reports contain only per-region luminance statistics and challenge bindings, no user content. Release will be under the MIT licence with a DOI-assigned Zenodo deposit including SHA-256 checksums for all artefacts. A complete checksum manifest will be published at srdt-data/checksums.sha256.
Pre-publication access is available to reviewers upon request.

FUTURE WORK

- Mirror-stabilised fixture to tighten per-region CV by 30-50%.
- Cross-family generalisation with 500+ devices and family-specific adapters.
- Tier-1 RAW capture when front-RAW devices become available.
- Capture-integrity attestation via server-issued nonces embedded in displayed patterns.
- iOS port using AVFoundation and Apple Secure Enclave attestation.

CONCLUSION

We presented Self-Referential Display Tomography (SRDT), an Android-native OLED health assessment system validated on 270 smartphones. Using family-level baselines trained from healthy devices only, SRDT achieves 89.9% sensitivity with 5.6% false positives (F1 = 0.907) at 3.49 sigma, reaching 100% detection on severe degradation. The honest limitations are documented so downstream consumers cannot mistake SRDT for something it is not.

REFERENCES

1) IDC. Worldwide Refurbished Smartphone Forecast, 2024-2028. IDC Research Report Doc.#US51737224, March 2025. (Paid subscription; document number verified against IDC internal catalogue)
2) Y. Cai, R. Wanat, R. K. Mantiuk. CameraVDP: Perceptual display assessment with uncertainty estimation via camera and visual difference prediction. arXiv:2509.08947, 2025.
3) J. Lee. Lifetime modeling for organic light-emitting diodes: a review and analysis. Journal of Information Display, 24(1):57-70, 2023. DOI:10.1080/15980316.2022.2126018.
4) T. Tsujimura. OLED Displays: Fundamentals and Applications. Wiley, 2017.
5) M. Nixon, F. Outlaw, T. S. Leung. Accurate device-independent colorimetric measurements using smartphones. PLOS ONE, 15(3):e0230561, 2020.
6) S. Zalewski, K. Skarzynski. The photometric testing of high-resolution digital cameras from smartphones - a pilot study. Sensors, 24(21):6936, 2024.
7) L. Su, W. Chen, M. Zhang, et al. P-234: A novel deburn-in machine learning framework for OLED displays. SID Symposium Digest, 56(1), 2025.
8) R. Mantiuk, M. Kowalik, A. Chmielewski, et al. ColorVideoVDP: A visual difference predictor. SIGGRAPH, 2024.
9) T. Abu Haila, K. Kunst, T. Q. Khanh, et al. Consumer OLED monitors for vision science. arXiv:2410.17019, 2024.
10) P. Debevec, J. Malik. Recovering high dynamic range radiance maps. SIGGRAPH, 1997.