Blockchain-Anchored Provenance Verification of Crop-Disease Images Using Perceptual Hashing and Metadata Consistency Checks

A Simulation Study for Agricultural Extension and Crop-Insurance Evidence

[1]Sonia Mohapatra, [2]Pravat Satpathy, [3]Saswat Swain
[1]Researcher, [2]Researcher, [3]Ph.D Scholar
[1,2,3]*Research and Development Division,*
[1,2,3]*Tekstein Solutions Private Limited, Bhubaneswar, Odisha, India*
[1]smctc98@gmail.com, analytics.pravat@gmail.com, saswat.mck@gmail.com, [2]analytics.pravat@gmail.com,
[3]saswat.mck@gmail.com

***Abstract*—Photographs of diseased leaves now serve as evidence: extension staff log them to follow outbreaks, and crop insurers accept them as proof of loss. Once such an image leaves the field, an ordinary file store cannot show whether its label, plot or date was changed, a lesion pasted in or painted out, or an old picture resubmitted. This paper describes AgriProv, a provenance scheme that binds content-aware image fingerprints to a permissioned ledger run by a small set of known validators. Each image is segmented, described by colour and texture statistics, and fingerprinted with a cryptographic hash plus global and tile-level perceptual hashes. Four consistency checks covering plot boundary, capture and upload time, label plausibility and near-duplicate search screen each submission before its digests, metadata hash and a hash-chained processing history are batched under a Merkle root. A Python implementation, mirrored in MATLAB, was tested on synthetic leaf images from 48 smallholder plots under nineteen scenarios. The scheme accepted 99.2 percent of benignly re-encoded images and every untouched one, rejected all metadata edits and image substitutions, and caught 95.3 percent of lesion splices covering two percent of the leaf, compared with 26.0 percent for a global perceptual hash alone. Registration checks flagged 95.3 percent of mislabelled and 98.7 percent of neighbouring-plot submissions, at a 5.4 percent false-flag rate driven mainly by late uploads. Each on-chain record needs 359 bytes, about 2.2 percent of the mean image. Three limitations are reported: a three percent crop defeats every perceptual test, the descriptor distance added no measurable detection, and clock rollback shorter than the upload window goes unnoticed.**

***Index Terms*—Agricultural data integrity, blockchain, crop insurance, image provenance, Merkle tree, perceptual hashing, plant disease, proof of authority, tamper detection.**

I. Introduction

A smartphone photograph of a blighted tomato leaf carries more than pixels. It implies a crop, a plot, a date and a diagnosis, and each of these may decide whether an advisory is issued for a block or whether a farmer is paid for a loss. Picture-based crop insurance has been shown to be workable precisely because such images are cheap to collect at scale [1], and national guidelines in India now provide for geotagged crop photographs captured through official mobile applications [2].

The weak point is the interval between capture and use. Field images travel through messaging services, district servers and contractor laptops, and at each hop the file may be recompressed, renamed or edited. A label can be switched from healthy to late blight; a picture taken on one plot can be filed under a neighbour's survey number; a lesion can be copied from another photograph to inflate severity, or erased to hide an outbreak. Label noise in plant-disease datasets is already documented [3], [4]; deliberate manipulation is less studied, does not average out, and occurs where money changes hands.

Two families of tools address parts of this problem. Blockchain traceability systems record supply-chain events in append-only ledgers [5]–[8], and smart contracts have been proposed for index-based crop insurance [9]. These systems record what a party *declares*; they do not examine the image. Image-forensic and perceptual-hashing methods do examine content [10]–[14], and recent work has anchored perceptual hashes of remote-sensing scenes [15] and videos [16] on blockchains. None of the reviewed work checks agronomic metadata, such as plot boundaries, capture time and disease label, against the image and against each other before the record is frozen. Once a false record is on a ledger, immutability only preserves the falsehood.

This paper treats provenance of a crop-disease image as two linked questions: was the record plausible when it was registered, and has anything changed since? The contributions are as follows.

1) A registration gate of four checks (geofence, capture and anchoring time, classifier label plausibility and tile-aware near-duplicate search) that screens submissions before they become immutable.
2) A compact on-chain record of 359 bytes combining SHA-256, a global and a 4×4 tile discrete cosine transform (DCT) perceptual hash, a half-precision colour and texture descriptor, a metadata hash and a hash-chained processing history, batched under Merkle roots in a proof-of-authority (PoA) ledger.
3) A verification rule that accepts benign re-encoding but localises small lesion splices, together with two analytical properties (tamper evidence and chance-match probability) checked against simulation.

4) A controlled evaluation of nineteen scenarios against three baselines that emulate exact-hash anchoring, global perceptual-hash anchoring and rule-only screening, with negative results reported in full.

II. Related Work

Deep convolutional networks reach high accuracy on curated leaf collections such as PlantVillage [17], and recent reviews document a wide range of architectures [18]. The same reviews and later analyses caution that laboratory images misrepresent field conditions and that label noise is common [3], [4].

Classical image forensics detects splicing and copy-move edits from statistical traces [12], and deep detectors now dominate the copy-move and splicing literature, though their behaviour under recompression and on unseen content remains uneven [13], [14]. Perceptual hashing takes the opposite stance: instead of searching for traces of editing, it compares an image with a stored fingerprint and tolerates content-preserving changes [10], [11].

On the ledger side, the hash-chained, timestamped block structure [19] and Merkle batching [20] give tamper evidence, and permissioned consensus such as PoA trades openness for low cost and fast finality [21], [22]. Agricultural deployments have focused on supply-chain traceability [5]–[8] and insurance automation [9]. Hasan and Salah anchored video provenance with exact hashes and smart contracts [23]; exact hashes, however, break under any re-encoding. Xu *et al.* combined perceptual hashing and blockchain for remote-sensing integrity [15], and Madapati and Pradhan did so for video copyright [16]. The C2PA specification defines signed provenance manifests bound to media at capture [24], which presumes capable capture devices. Table I positions the present work.

TABLE I
Comparison With Related Work

Ref.	Object	Integrity primitive	Content test	Field checks
[4], [17], [18]	Leaf images	None	Classifier	No
[5]–[8]	Agri-food lots	Ledger + IoT	No	Partial
[9]	Index insurance	Smart contract	No	Weather index
[23]	Video	Exact hash	No	No
[15]	Remote sensing	Ledger + pHash	Global	No
[16]	Video	Ledger + pHash	Global	No
[24]	General media	Signed manifest	Hard binding	Capture claims
[1], [2]	Crop photos	Central store	Manual	Geotag
This work	Crop-disease image	PoA ledger, SHA-256, pHash	Global + tiles	C1–C4

III. Problem Formulation and Threat Model

A *submitter* (farmer, surveyor or extension worker) captures an image I with a field application that reads a satellite positioning fix x_{dev} and the device clock t_{dev}. The submitter adds claimed metadata m = (plot f, label y, capture time t_c, position, crop, season). A *verifier* (insurer, auditor or extension officer) later receives a package (I', m', record identifier) and must decide whether it matches what was registered. A farm registry gives each plot a centre c_f and radius r_f.

Two threat classes are considered. **Post-registration tampering** (T1–T6) alters a registered package: label swap, plot and position swap, timestamp shift, lesion splice, lesion erasure, or whole-image substitution. **Registration-time fraud** (R1–R4) submits a false package in the first place: mislabelling, misattribution to a neighbouring or distant plot, backdating the claimed date or the device clock, and reusing an image already on record. Benign processing (re-encoding, resizing, brightness change) must not be flagged.

The assumptions are stated explicitly. A1: SHA-256 is collision resistant and Ed25519 signatures are unforgeable. A2: fewer than $q = \lfloor 2V/3 \rfloor + 1$ of the V validators collude. A3: the device position and clock are honest unless the attacker is modelled as rolling the clock back (R3b); a fully forged sensor stream is out of scope and is discussed in Section VIII. A4: the registry of plot boundaries is correct. The goal is high acceptance of honest packages and high detection of attacks at district-scale cost.

IV. Proposed Framework

A. Architecture

Fig. 1 shows the data flow. The field application captures the image and metadata; an analysis stage computes masks, a descriptor and fingerprints; the registration gate raises flags; a record builder signs a compact record; and validators seal records into blocks. Image bytes are kept off-chain in a store addressed by their SHA-256 digest, so the ledger carries only fixed-size records.

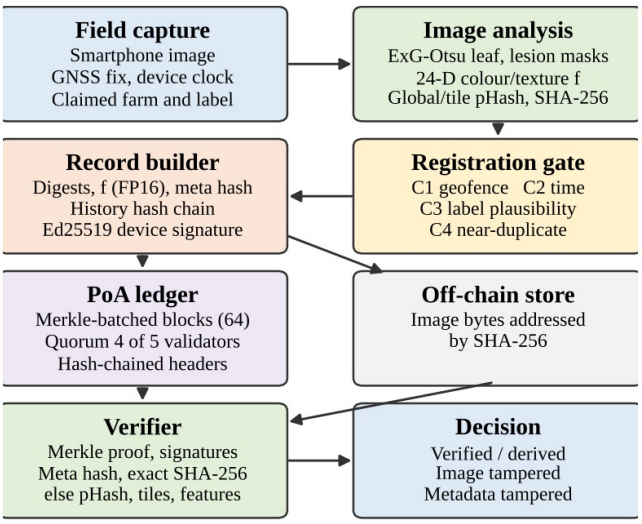


Fig. 1. AgriProv architecture. Registration runs from capture to the ledger, with gate flags carried into the record; the verifier fetches the record and its Merkle proof and returns one of three decisions.

B. Leaf Segmentation and Descriptor

With chromatic coordinates $r = R/(R+G+B)$, and g, b defined likewise, the excess-green index is

(1)

and the leaf mask is the set of pixels with ExG above the Otsu threshold [25],

(2)

where ω_0, ω_1 are the class probabilities and μ_0, μ_1 the class means of the two sides of t. An opening with a disc of radius 2, removal of components under 500 pixels, a closing and hole filling give the leaf region Ω. Lesions appear as non-green pixels inside Ω, so the lesion mask is

(3)

with t_ρ the Otsu threshold of ρ over Ω; the cap of 0.52 prevents a healthy leaf from being split in two. The descriptor $\mathbf{f} \in \mathbb{R}^{24}$ holds the mean and standard deviation of hue, saturation and value over Ω, an 8-bin hue histogram, lesion fraction $|\Lambda|/|\Omega|$, lesion count and mean size, mean lesion colour, and four grey-level co-occurrence (GLCM) statistics [26] averaged over four directions at unit distance. For a normalised GLCM $P(i, j)$, contrast and homogeneity are

(4)

(5)

C. Perceptual and Cryptographic Fingerprints

The grey image is resized to 32×32 pixels and transformed with a two-dimensional DCT; the 64 lowest-frequency coefficients D_k form the hash [10], [11], with $[\cdot]$ the Iverson bracket:

(6)

(7)

A global hash $\mathbf{h}_g$ is computed on the whole image and sixteen tile hashes $\mathbf{h}_1, \ldots, \mathbf{h}_{16}$ on a 4×4 grid of 64×64 tiles. The global hash tolerates re-encoding; the tile hashes confine a local edit to the one or two tiles it touches, where it changes a larger share of the coefficients. SHA-256 of the encoded bytes, $\sigma = H(I)$, gives an exact-match fast path.

D. Registration Gate

Four checks run before a record is sealed. A flag does not block registration; it marks the record for review, which keeps evidence of attempted fraud on the ledger. The plot check flags a submission whose position falls outside the claimed plot by more than a tolerance δ:

(8)

The time check compares the claimed date with the device clock and bounds the delay between capture and anchoring time t_a:

(9)

The label check uses a multinomial logistic classifier on standardised **f** and flags a claimed label y whose posterior is below θ; the near-duplicate check searches the index of registered hashes:

(10)

(11)

where d_t^j is the largest of the sixteen tile distances to record j. Requiring both conditions matters: two different leaves photographed against similar soil can share a global hash more easily than sixteen tile hashes.

E. Ledger Record and Blocks

Each record stores identifier, σ, $\mathbf{h}_g$, the sixteen tile hashes, **f** in half precision, the metadata hash $\mu = H$(canonical JSON of m), label, plot, capture time and a processing-history digest built as a hash chain over the analysis steps (segmentation, features, classifier) with their versions:

(12)

The submitting device signs the record with Ed25519. Validators batch $B = 64$ records, compute a Merkle root with domain-separated leaf and node hashes, $H(0x00 \| \text{rec})$ and $H(0x01 \| L \| R)$ [20], and chain block headers by hash [19]. A block is valid when q of $V = 5$ validators, chosen round-robin from the proposer, sign its header; with $V = 5$ the quorum is $q = 4$. PoA suits this setting because validators are known institutions such as an agriculture department, an insurer and a university, and because energy and latency are low [21], [22].

F. Verification Rule and Algorithm

Given a package (I', m') and its record r, the verifier checks the Merkle proof and header signatures, then accepts if and only if

(13)

(14)

where Π is the perceptual test, d_g the global Hamming distance, d_t the largest tile distance and $d_f = \|(\mathbf{f}' - \mathbf{f}) \oslash \mathbf{s}\|_2 / \sqrt{24}$ the descriptor distance scaled by the training standard deviations **s**. A failed first term means metadata tampering; a failed second term means image tampering. Thresholds are the 99th percentiles of each distance over benign transforms of a separate calibration set. Algorithm 1 summarises both phases.

Algorithm 1 Registration and verification

Register(I, m, x_{dev}, t_{dev}):
Ω, Λ ← Segment(I); **f** ← Descriptor(I, Ω, Λ)
σ ← SHA-256(I); $\mathbf{h}_g$, {$\mathbf{h}_k$} ← pHash(I)
flags ← {C1, C2, C3, C4} (Eqs. 8–11)
η ← HistoryChain(σ, steps); μ ← $H(m)$
rec ← Sign_{dev}(σ, **h**, **f**, μ, η, flags)
append rec to pending; if |pending| = B: seal block
Verify(I', m', id):
rec, proof ← Ledger(id); check proof and q signatures
if $H(m') \neq$ rec.μ: return METADATA TAMPERED
if SHA-256(I') = rec.σ: return VERIFIED
compute d_g, d_t, d_f; if all within τ: return DERIVED
return IMAGE TAMPERED

Complexity. For an image of n pixels, segmentation, descriptor and hashing cost $O(n)$. The C4 search is a linear scan, $O(N)$ in the number N of registered images. A Merkle proof has $\lceil \log_2 B \rceil = 6$ siblings, so verification of the ledger part is $O(\log B)$ and independent of chain length. Storage grows as $O(N)$.

V. Analytical Properties

Proposition 1 (tamper evidence). Under A1 and A2, any change to a committed record or block header makes chain verification fail. *Proof sketch:* a changed record changes its leaf hash and, by collision resistance, the Merkle root; a changed root changes the header, which invalidates the q validator signatures and breaks the prev-hash link of the next block. A2 excludes re-signing. □

Proposition 2 (chance match). If the 64 bits of hashes from unrelated images are independent and equiprobable, the probability that two unrelated images fall within distance τ is

(15)

For the calibrated $\tau_g = 2$, Eq. 15 gives 1.13×10^{-16}. Independence is an idealisation; Section VII-E tests it.

Storage. With record size s_r and block overhead s_b (header plus q signatures), the on-chain volume after N images is

(16)

Here $s_r = 359$ B (including the 64-B device signature) and $s_b = 342$ B, so each image costs 364.3 B on-chain.

VI. Simulation Setup

Data. No public crop-disease collection carries trustworthy plot, time and edit history, which is the very gap studied here, and attacks need pixel-level ground truth. A procedural generator was therefore built. It draws a tapered elliptical leaf

with midrib and texture on a soil background and paints four classes: healthy (35 percent carry abiotic chlorosis patches as a confuser), leaf spot (dark spots with yellow halos), early blight (large irregular lesions with concentric rings) and rust (small orange pustules). Severity is drawn from 0.5–3 percent for one fifth of images and 3–18 percent otherwise; illumination gain and sensor noise vary per image. Images are 256×256 pixels at JPEG quality 92. The farm registry contains six villages of eight adjacent plots (2×4 grid, radius 55–110 m, 4-m gaps). Positions are uniform within the plot plus a 6-m positioning error, or 25 m with probability 0.05; anchoring delay is exponential with a 10-h mean, or 72–168 h with probability 0.05 to model offline areas.

Implementation. The reference implementation uses Python 3.13 with NumPy, SciPy, scikit-image, scikit-learn and the *cryptography* library, with global seed 2026. A MATLAB script reproducing the same pipeline with the Image Processing and Statistics toolboxes is released with the code; its signatures use HMAC instead of Ed25519, and all figures in this paper come from the Python run. Timing was measured on one core of a 2.1-GHz Intel Xeon virtual machine with 8 GB of memory. Table II lists the parameters.

TABLE II
SIMULATION PARAMETERS

Parameter	Value
Training / calibration / registered images	800 / 240 / 800
Test packages per scenario	300 (erasure: diseased only)
Plots; validators; quorum; block size	48; 5; 4; 64
τ_g, τ_t, τ_f (99th pct. of benign)	2 bits, 4 bits, 2.27
δ, Δ_c, Δ_a, θ	10 m, 24 h, 72 h, 0.15
Benign set (calibrated)	JPEG 75; resize to 0.75; brightness 1.10
Out-of-envelope benign	3 percent crop per side + rescale

Baselines. Three baselines emulate the main prior designs on identical data; they are emulations, not reimplementations of the cited systems. BL1 anchors exact SHA-256 hashes and metadata hashes, as in [23]. BL2 anchors a global DCT perceptual hash with the same τ_g, as in [15], [16], and deduplicates with it. BL3 applies the content and metadata rules without a ledger, as a central screening service would; in post-registration attacks its adversary edits metadata self-consistently. All four methods share the same classifier and thresholds.

Metrics. The provenance verification rate (PVR) is the fraction of genuine or benign packages accepted. The tamper-detection rate (TDR) is the fraction of attacked packages rejected (post-registration) or flagged (registration time). The false-flag rate (FFR) is the fraction of honest registrations flagged by any check. Processing time per stage and bytes per record complete the set.

VII. RESULTS AND DISCUSSION

A. Classifier and False Flags on Honest Data

The classifier reached 98.1 percent accuracy on the 800 registered images. Nearly all errors were early-blight leaves read as healthy (11 of 200), mostly low-severity cases whose few lesions resemble chlorosis. This figure is optimistic for synthetic data [3], [4]. On honest registrations the gate flagged 5.38 percent: C1 0.88, C2 4.63, C3 0.13 and C4 0.00 percent. C2 dominates because of the simulated offline uploads; widening Δ_a lowers this rate but widens the rollback window examined in Section VII-D.

B. Post-Registration Tampering

TABLE III
POST-REGISTRATION RESULTS (PERCENT)

Scenario	BL1	BL2	BL3	Prop.
PVR, higher is better				
Genuine, unchanged	100.0	100.0	99.3	100.0
B1 JPEG 75	0.0	99.0	98.3	99.0
B2 resize 0.75	0.0	100.0	78.7	99.7
B3 brightness 1.10	0.0	99.0	70.3	99.0
B4 crop 3 % (outside)	0.0	0.3	90.3	0.0
TDR, higher is better				
T1 label swap	100.0	100.0	97.0	100.0
T2 plot + position	100.0	100.0	0.7	100.0
T3 time shift	100.0	100.0	0.7	100.0
T4 lesion splice 2 %	100.0	26.0	21.3	95.3
T5 lesion erasure	100.0	33.3	63.4	99.5
T6 substitution	100.0	100.0	1.0	100.0

n = 300 per row except T5 (n = 186, diseased leaves only).

Table III summarises the post-registration results. Exact hashing (BL1) detects every attack but accepts no benign variant: a single pass through a messaging application would make an honest farmer's evidence fail. Rule-only screening (BL3) misses self-consistent metadata edits (T2, T3) and substitutions almost entirely, because nothing records what was originally submitted; it also rejects 21.3 and 29.7 percent of resized and brightened images, since the classifier posterior drifts with global colour changes. Content plausibility is therefore better applied once, at registration, on the original upload.

The global-hash baseline (BL2) handles metadata edits and substitution, but detects only 26.0 percent of 2-percent splices and 33.3 percent of erasures. The proposed rule raises these to 95.3 and 99.5 percent at the same benign acceptance (99.2 against 99.3 percent). Fig. 2 shows why: a splice on one tile moves that tile's hash by 12 bits while the global hash moves by 2.

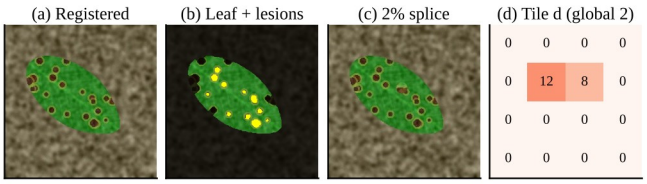


Fig. 2. (a) Registered leaf-spot image; (b) leaf region (bright) and lesion mask (yellow); (c) the same image after a 2-percent lesion splice; (d) tile-wise Hamming distances between (a) and (c).

The 3-percent crop (B4) is a clear negative result. All perceptual tests reject it, because cropping shifts the DCT grid and every tile boundary. BL3 accepts it only because it never compares images. Cropping is common when users frame a leaf, so the field application should either forbid it after capture or log it as a declared processing step in the history chain.

C. Splice Size and Ablation

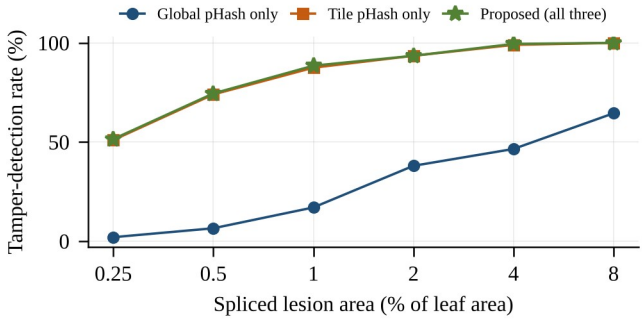


Fig. 3. Tamper-detection rate against spliced lesion area (200 images per point). Global + descriptor coincides with global only, and global + tiles with the proposed rule (Table IV).

Fig. 3 sweeps the spliced area. Tile hashes detect 51 percent of splices covering 0.25 percent of the leaf and 87.5 percent at 1 percent, while the global hash stays at 2 and 17 percent. Table IV isolates each component.

TABLE IV
ABLATION OF THE IMAGE TEST (PERCENT)

Test	Benign PVR	Splice 1 %	Splice 2 %	Erasure
Global only	99.3	17.0	38.0	35.2
Global + descriptor	99.3	17.0	38.0	35.2
Tiles only	99.8	87.5	93.5	100.0
Global + tiles	99.2	88.5	93.5	100.0
All three	99.2	88.5	93.5	100.0

Erasure on 200 sampled records, diseased leaves only; substitution was detected at 100 percent by every perceptual variant.

The descriptor distance contributes nothing measurable. Its calibrated threshold, $\tau_f = 2.27$, is large because brightness and resizing shift hue, value and texture statistics by about as much as a small splice does. This was not the expected outcome; the descriptor is kept in the record because C3 uses it and because it gives auditors a semantic snapshot of the registered image, but it should not be counted as a detection layer. Tiles carry almost all the gain.

D. Registration-Time Fraud

TABLE V
REGISTRATION-TIME RESULTS (PERCENT)

Scenario	Prop.	Main check	BL1	BL2
R0 honest (1 − FFR)	92.7	–	100.0	100.0
R1 mislabel	95.3	95.3	0.0	0.0
R2a adjacent plot	98.7	98.7	0.0	0.0
R2b distant plot	100.0	100.0	0.0	0.0
R3a backdated claim	100.0	100.0	0.0	0.0
R3b clock rollback	86.0	85.3	0.0	0.0
R4a reuse, same file	100.0	100.0	100.0	100.0
R4b reuse, re-encoded	98.3	98.3	0.0	98.3

Prop. counts a flag from any check and equals BL3, which applies the same gate; the main-check column isolates the check aimed at each attack.

Table V shows that the gate, not the ledger, carries registration-time defence; ledger-only baselines detect nothing except exact or near duplicates. Misattribution to a neighbouring plot is caught at 98.7 percent by C1; the misses are photographs taken within δ of a shared boundary. Re-encoded reuse escapes exact hashing but is caught at 98.3 percent by C4, with no false duplicates among honest submissions. Clock rollback is detected only when the rollback plus upload delay exceeds $\Delta_a = 72$ h (85.3 percent of the 1–10-day rollbacks simulated); shorter rollbacks are invisible without a trusted time source.

E. Validity of the Independence Model

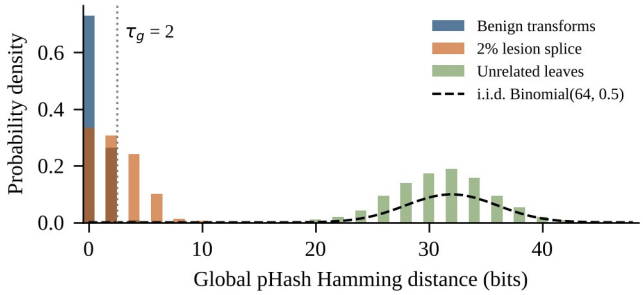


Fig. 4. Global hash distances for benign transforms, 2-percent splices and unrelated leaf pairs, with the independent-bit Binomial(64, 0.5) model of Proposition 2.

Fig. 4 compares the distance distributions. Benign transforms stay within $\tau_g = 2$ bits in 99 percent of calibration cases (maximum 4). Splices spread from 0 to 10 bits, with 64 percent at or below τ_g, which is why the global hash alone misses most of them. Unrelated pairs centre on 31.1 bits (standard deviation 4.25) against 32 and 4 for the independent model. The smallest of 3998 unrelated distances was 16 bits, so the empirical chance-match rate at τ_g is zero, consistent with Eq. 15. The lower tail is nonetheless heavier than the model predicts: 1.15 percent of pairs lie within 20 bits against 0.18 percent, a 6.2-fold excess, so Eq. 15 understates chance matches at looser thresholds. The expected collapse of hash spacing for uniform leaf imagery was therefore only partly borne out, and the margin at τ_g remains wide; but real photographs with repeated backgrounds may behave differently.

F. Processing Time and Storage

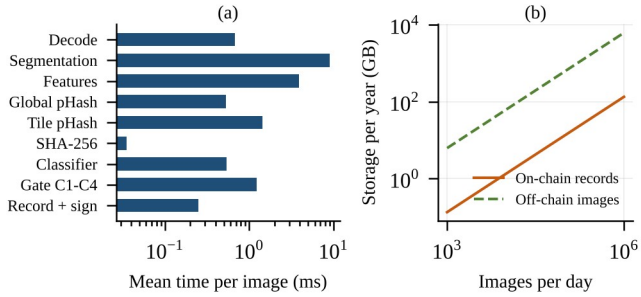


Fig. 5. (a) Mean registration time per stage; (b) projected yearly storage on-chain and off-chain against daily image volume.

Registration took 17.6 ms per image, of which segmentation took 9.0 ms; tile hashing added 1.4 ms (Fig. 5a). Verification of an unchanged package (Merkle proof, signatures, metadata and exact hash) took 0.26 ms; the perceptual path for a re-encoded image took 11.8 ms. Sealing a 64-record block took 0.26 ms, and verifying the full 13-block chain took 5.9 ms. In 200 trials of record modification and one attempt to recompute a Merkle root without validator keys, chain verification failed every time, as Proposition 1 predicts. The on-chain cost is 364.3 B per image, 2.2 percent of the mean 16.9-kB JPEG; at 10^5 images per day this is 13.3 GB per year on-chain against 617 GB off-chain (Fig. 5b).

VIII. LIMITATIONS

1) All images are synthetic, so segmentation and classification are optimistic; occlusion, shadow and mixed infections in

field photographs will degrade both. The descriptor-distance test gave no gain at its calibrated threshold.

2) Cropping, rotation and perspective change defeat the DCT hashes; keypoint-based hashes [10] or declared-crop logging should be evaluated.
3) Device position and clock are trusted. A spoofed fix or a rollback shorter than Δ_a passes the gate; capture-time signing with trusted timestamps [24] would close part of this gap.
4) Duplicate search is linear; the ledger ran in one process without network delay or the consistency issues analysed in [21]; the MATLAB script was not used for the reported numbers.

IX. Conclusion and Future Work

AgriProv anchors crop-disease images to a permissioned ledger through a compact record that combines exact and perceptual fingerprints with a processing history, and screens each submission with plot, time, label and duplicate checks before it becomes immutable. In simulation, the scheme kept benign re-encoded images verifiable, detected every metadata edit and substitution, and raised detection of 2-percent lesion splices from 26.0 to 95.3 percent over a global-hash ledger, at 364 B and about 18 ms per image. The results also show where the design stops: crops break the hashes, the descriptor test added nothing, and short clock rollbacks need trusted time. Future work will test the pipeline on field images collected with a signed capture application, replace linear duplicate search with an indexed structure, and run the ledger across independent validator nodes.

Acknowledgment

The authors thank colleagues in the Research and Development Division of Tekstein Solutions Private Limited for discussions on field-data workflows. Code, synthetic data and result logs are available from the corresponding author.

References

[1] F. Ceballos, B. Kramer, and M. Robles, “The feasibility of picture-based insurance (PBI): Smartphone pictures for affordable crop insurance,” *Dev. Eng.*, vol. 4, Art. no. 100042, 2019, doi: 10.1016/j.deveng.2019.100042.

[2] Dept. Agriculture & Farmers Welfare, Govt. of India, “Pradhan Mantri Fasal Bima Yojana (PMFBY): Operational guidelines,” New Delhi, India, 2023.

[3] J. G. A. Barbedo, “Deep learning applied to plant pathology: The problem of data representativeness,” *Trop. Plant Pathol.*, vol. 47, no. 1, pp. 85–94, Feb. 2022, doi: 10.1007/s40858-021-00459-9.

[4] M. Xu, H. Kim, J. Yang, A. Fuentes, Y. Meng, S. Yoon, T. Kim, and D. S. Park, “Embracing limited and imperfect training datasets: Opportunities and challenges in plant disease recognition using deep learning,” *Front. Plant Sci.*, vol. 14, Art. no. 1225409, Sep. 2023, doi: 10.3389/fpls.2023.1225409.

[5] A. Kamilaris, A. Fonts, and F. X. Prenafeta-Boldú, “The rise of blockchain technology in agriculture and food supply chains,” *Trends Food Sci. Technol.*, vol. 91, pp. 640–652, Sep. 2019, doi: 10.1016/j.tifs.2019.07.034.

[6] S. A. Bhat, N.-F. Huang, I. B. Sofi, and M. Sultan, “Agriculture-food supply chain management based on blockchain and IoT: A narrative on enterprise blockchain interoperability,” *Agriculture*, vol. 12, no. 1, Art. no. 40, Jan. 2022, doi: 10.3390/agriculture12010040.

[7] R. M. Ellahi, L. C. Wood, and A. E.-D. A. Bekhit, “Blockchain-based frameworks for food traceability: A systematic review,” *Foods*, vol. 12, no. 16, Art. no. 3026, Aug. 2023, doi: 10.3390/foods12163026.

[8] B. J. Ordóñez García, A. Alexopoulos, K. Koutras, A. Kalogeras, K. Stefanidis, and V. M. Martos Núñez, “Blockchain in agriculture: A PESTELS analysis,” *IEEE Access*, vol. 11, pp. 73647–73679, 2023, doi: 10.1109/ACCESS.2023.3295889.

[9] I. A. Omar, R. Jayaraman, K. Salah, H. R. Hasan, J. Antony, and M. Omar, “Blockchain-based approach for crop index insurance in agricultural supply chain,” *IEEE Access*, vol. 11, pp. 118660–118675, 2023, doi: 10.1109/ACCESS.2023.3327286.

[10] V. Monga and B. L. Evans, “Perceptual image hashing via feature points: Performance evaluation and tradeoffs,” *IEEE Trans. Image Process.*, vol. 15, no. 11, pp. 3452–3465, Nov. 2006, doi: 10.1109/TIP.2006.881948.

[11] Y. Zhou, X. Li, C. Xiong, H. Yao, and C. Qin, “A survey of perceptual hashing for multimedia,” *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 21, no. 7, pp. 1–28, Jul. 2025, doi: 10.1145/3727880.

[12] H. Farid, “Image forgery detection,” *IEEE Signal Process. Mag.*, vol. 26, no. 2, pp. 16–25, Mar. 2009, doi: 10.1109/MSP.2008.931079.

[13] M. Zanardelli, F. Guerrini, R. Leonardi, and N. Adami, “Image forgery detection: A survey of recent deep-learning approaches,” *Multimedia Tools Appl.*, vol. 82, no. 12, pp. 17521–17566, May 2023, doi: 10.1007/s11042-022-13797-w.

[14] I. Shallal, L. Rzouga Haddada, and N. Essoukri Ben Amara, “Image forgery detection with focus on copy-move: An overview, real world challenges and future directions,” *Appl. Sci.*, vol. 15, no. 21, Art. no. 11774, Nov. 2025, doi: 10.3390/app152111774.

[15] D. Xu, N. Ren, and C. Zhu, “Integrity authentication based on blockchain and perceptual hash for remote-sensing imagery,” *Remote Sens.*, vol. 15, no. 19, Art. no. 4860, Oct. 2023, doi: 10.3390/rs15194860.

[16] S. L. Madapati and N. R. Pradhan, “Decentralizing video copyright protection: A novel blockchain-enabled framework with performance evaluation,” *Front. Artif. Intell.*, vol. 8, Art. no. 1655709, Aug. 2025, doi: 10.3389/frai.2025.1655709.

[17] S. P. Mohanty, D. P. Hughes, and M. Salathé, “Using deep learning for image-based plant disease detection,” *Front. Plant Sci.*, vol. 7, Art. no. 1419, Sep. 2016, doi: 10.3389/fpls.2016.01419.

[18] M. Shoaib, B. Shah, S. El-Sappagh, A. Ali, A. Ullah, F. Alenezi, T. Gechev, T. Hussain, and F. Ali, “An advanced deep learning models-based plant disease detection: A review of recent research,” *Front. Plant Sci.*, vol. 14, Art. no. 1158933, Mar. 2023, doi: 10.3389/fpls.2023.1158933.

[19] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system,” 2008. [Online]. Available: https://bitcoin.org/bitcoin.pdf

[20] R. C. Merkle, “A digital signature based on a conventional encryption function,” in *Advances in Cryptology—CRYPTO ’87* (Lecture Notes in Computer Science, vol. 293), C. Pomerance, Ed. Berlin, Germany: Springer, 1988, pp. 369–378, doi: 10.1007/3-540-48184-2_32.

[21] S. De Angelis, L. Aniello, R. Baldoni, F. Lombardi, A. Margheri, and V. Sassone, “PBFT vs proof-of-authority: Applying the CAP theorem to permissioned blockchain,” in *Proc. 2nd Italian Conf. Cyber Security (ITASEC)*, Milan, Italy, Feb. 2018, CEUR Workshop Proc., vol. 2058, paper 6.

[22] J. Xu, C. Wang, and X. Jia, “A survey of blockchain consensus protocols,” *ACM Comput. Surv.*, vol. 55, no. 13s, pp. 1–35, 2023, doi: 10.1145/3579845.

[23] H. R. Hasan and K. Salah, “Combating deepfake videos using blockchain and smart contracts,” *IEEE Access*, vol. 7, pp. 41596–41606, 2019, doi: 10.1109/ACCESS.2019.2905689.

[24] Coalition for Content Provenance and Authenticity (C2PA), “Content Credentials: C2PA technical specification,” ver. 2.3, 2026. [Online]. Available: https://spec.c2pa.org/specifications/specifications/2.3/specs/_attachments/C2PA_Specification.pdf

[25] N. Otsu, “A threshold selection method from gray-level histograms,” *IEEE Trans. Syst., Man, Cybern.*, vol. 9, no. 1, pp. 62–66, Jan. 1979, doi: 10.1109/TSMC.1979.4310076.

[26] R. M. Haralick, K. Shanmugam, and I. Dinstein, “Textural features for image classification,” *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-3, no. 6, pp. 610–621, Nov. 1973, doi: 10.1109/TSMC.1973.4309314.