

essentialism: An On-Chain Generative Work That Forgets Two Photographs

Benjamin Meadows
Independent Artist and Researcher
ben@benjaminmeadows.com

October 2026

Abstract

essentialism is a generative artwork of 44 editions whose program is stored on the Tezos blockchain. It carries two of my photographs inside its code as 28×38 grids of 16 gray levels. Each edition samples its grid at a weight fixed by the edition number, 100% at edition 1 falling linearly to 12% at edition 44, mixed with value noise whose weight rises as the photograph's falls, and each token states that weight in its trait line as its "memory". The ledger keeps the program, and in its block history every superseded version of it; the editions are where the work forgets. Over 1,000 seeds and all 44 editions (44,000 renders) I measured how strongly each photograph still shows in the generator's internal darkness field and in the ink of the drawn dots. The field keeps a median correlation of at least 0.94 with the photograph through edition 21 and loses 0.43 (tree) and 0.44 (sky) of it between editions 34 and 44: a signal weight that falls linearly against a rising noise weight gives a correlation that holds and then drops. A closed form with two measured variances reproduces the field curve to within 0.010 for one photograph and 0.021 for the other. Counting every render, the dots' ink keeps a median rank correlation of 0.82 with the tree photograph at edition 1 and 0.14 at edition 44, against medians between 0.00 and 0.104 for a counterfactual null in which the other photograph drives a re-render of the same seed. The decline is orderly in the medians, not token by token, and the memory percentage printed on each token is above the median drawn correlation through edition 11 and below it from edition 12, with one exception. Nine editions are minted; the latest states a memory of 84%, and renders at its edition number have a median drawn correlation of 0.79 (tree) or 0.81 (sky). The ledger also kept something I would have chosen to forget: the first published version of the code claimed an algorithm that was not mine. Separately, the platform's image of a minted token could not be reproduced from the seed data it stores on chain. The generator, measurement scripts, data and this paper's source rebuild from one command; they are archived at <https://doi.org/10.5281/zenodo.23333266> and maintained at <https://github.com/TheBenMeadows/essentialism-preprint>.

1 The work

essentialism [Meadows, 2026b] was published on 31 July 2026 as generator 407 on bootloader [objkt.com, 2026], a platform on Tezos [Goodman, 2014] that stores a generative program in a contract's storage and renders each token in the browser from that program and a per-token seed. Generators that live entirely on a ledger go back at least to Larva Labs' Autoglyphs, whose drawing code sits in its Ethereum contract [Larva Labs, 2019]. The edition is 44 tokens at 8 tez. Each token




Figure 1: One seed (1391882777, edition 1’s stored entropy reduced modulo 2^{32}) drawn at four edition numbers by the on-chain generator, version 6. Traits: field density, two inks, source “blue tree”, one ember. The dark mass at edition 1 is the trunk and crown of the tree. This is not the platform’s image of edition 1, which comes from a seed I could not derive (Section 5). Only the memory factor changes across the row, but because the dot sizes and run lengths depend on the memory field, the composition shifts as well. Rendered by `repro/figures.mjs`, rasterized with Inkscape 1.4.4.

is a 750×1000 SVG: columns of dots dropped down a sheet of paper, colored from a small library of inks.

The column structure is not mine. In November 2022 I worked through Andy Haskell’s tutorial on converting images to mosaics in `p5.js` [Haskell, 2022], randomized the dot size, pointed it at my own photographs, and saved nine test outputs under the name Essentialism without releasing them. In the tutorial’s version, and in mine, each dot asks a photograph for its color. The bootloader runtime allows no external files, so the on-chain program cannot load a photograph, and in July 2026 I had to decide what the dots would ask instead.

They ask two photographs that were reduced until they fit in the code. “Blue Tree” [Meadows, 2023a] is a photograph of a tree that was removed after a storm in December 2023, three months after I minted the image on Ethereum, so the picture has outlived its subject. “Natural Gradient” [Meadows, 2023b] is an unedited sky whose full-resolution file sits on Arweave [Williams et al., 2019]. Both originals are minted; how long their image files stay retrievable depends on the services that hold them. The work embeds a coarse copy of each and lets that copy fade across the edition: edition 1 samples its photograph at full strength, and edition 44 at 12%. The dots take their ink from what is left. Figure 1 shows one seed at four editions.

Every token states a “memory” percentage in its trait line. Before writing this paper I had not checked how much of a photograph that percentage leaves in the drawing.

2 Two ways of keeping

In relearning experiments, most of what is lost is lost early. Ebbinghaus measured his own relearning savings for nonsense syllables at intervals from twenty minutes to a month and found the steepest decline in the first hour [Ebbinghaus, 1885]; a replication with the same method reproduced the

curve [Murre and Dros, 2015], and in three experiments and a reanalysis of Ebbinghaus’s data, forgetting was best described by a power function of time, whose rate of loss falls with time [Wixted and Ebbesen, 1991]. Forgetting in this literature is a property of the rememberer.

Digital storage is often assumed to remove forgetting, and it does not. Bits decay and the systems holding them fail [Rosenthal, 2010]; one in five science, technology and medicine articles suffers from reference rot, and seven in ten of those that cite web resources do [Klein et al., 2014], and half of the URLs in United States Supreme Court opinions no longer pointed to the cited material when Zittrain, Albert and Lessig checked [Zittrain et al., 2014]. Chun describes digital media as an “enduring ephemeral”: always available in principle, regenerated constantly in practice, and liable to vanish [Chun, 2008]. Ernst argues that the digital archive is an operation rather than a store [Ernst, 2013], and Kirschenbaum shows how much of a file’s history persists physically even after it is “deleted” [Kirschenbaum, 2008]. Rinehart and Ippolito argue that the vulnerability of new media art, whose software and hardware change faster than institutions’ preservation methods, illustrates a larger crisis for social memory [Rinehart and Ippolito, 2014].

A blockchain promises the opposite for the bytes written to it. A ledger replicated across independent nodes and extended only by appending blocks [Nakamoto, 2008, Wood, 2014, Goodman, 2014] offers no operation that removes a past state; a contract can overwrite the current value of a key, but every earlier value stays in the block that wrote it, and archive nodes and indexers continue to serve it. Not every node does: in rolling mode a node keeps only recent blocks, and in full mode it keeps every block, including the transactions that carried each version’s code, but prunes old state and receipts [Octez developers, 2026]. The current value of a key is part of the live state every node holds; a superseded value survives because someone chooses to keep the history. The cost is paid once, per byte: on Tezos, 250 mutez (0.00025 tez) for every byte added to a contract’s storage. Superseded values are not charged again. For larger files, content addressing guarantees that a file is the one named, though not that anyone still serves it [Benet, 2014], and Arweave’s storage endowment is designed to pay for continued availability [Williams et al., 2019]. This permanence collides with law that treats forgetting as a right. Article 17 of the GDPR gives a data subject the right to erasure [European Parliament and Council of the European Union, 2016]; two years earlier, under the Data Protection Directive it replaced, the Court of Justice had required a search engine to delist links to outdated information about a person in *Google Spain* [Court of Justice of the European Union, 2014]. Finck sets out why an append-only ledger cannot comply in any straightforward way [Finck, 2018]. Mayer-Schönberger argues for building expiry into information systems by default [Mayer-Schönberger, 2009], and Bannon for treating forgetting as a design feature rather than a defect [Bannon, 2006]. Borges’s Funes, who cannot forget anything, also cannot think, because thinking requires discarding differences [Borges, 1944].

Artists have staged loss as a process for decades. Lucier rerecorded his voice in a room until the room’s resonances replaced the words [Lucier, 1970]; Basinski played decaying tape loops until the oxide shed from them [Basinski, 2002]; Metzger’s manifesto proposed works whose making was their destruction [Metzger, 1959]. Opalka’s count is the nearest to a fading sequence: from 1972 each canvas was painted on a ground about 1% whiter than the one before, so the series whitens as it goes on [Opalka, 1965]. Simon’s *Every Icon* runs a process toward an end no viewer will see [Simon, 1997]. On ledgers, Friend’s *Lifeforms* are tokens that die unless their holder gives them away within 90 days [Friend, 2021], and some editions already key the work to mint order: in Reas’s *923 EMPTY ROOMS* the mint number selects the composition [Reas, 2023], and in Gerrard’s *World Flag* release order follows each country’s greenhouse-gas ranking [Gerrard, 2023]. Weaker later copies are old in print: a drypoint plate’s burr wears away so quickly that only about a dozen good impressions can be taken [Department of Drawings and Prints, 2003], and Burn’s




Figure 2: The two photographs as the ledger holds them: 28×38 cells, 16 gray levels, read from the generator’s string literals by `repro/field.mjs`. This is the most of either photograph any edition can recall.

Xerox Book is a hundred successive photocopies of a blank sheet, bound in order [Burn, 1968]. On a ledger nothing wears, and I have not found an on-chain edition that schedules such a decline. Lucier’s tape, Basinski’s loops and Metzger’s objects lose their material. **essentialism** cannot: until the generator is updated again, its current program, and with it both grids, is part of the ledger’s live state, held by every node. The only thing the work can vary is the edition counter, so that is where the forgetting goes. Each token’s memory weight is fixed at mint and never changes afterwards; the schedule runs in the order the edition is minted, and earlier tokens keep what they were given. Barthes located a photograph’s force in its “that-has-been” [Barthes, 1981], a claim that the thing was there. Here the photograph stays in the code, and what the edition number lowers is how much of it the drawing shows.

3 Construction

The generator is 10,098 bytes of JavaScript (generator version 6, SHA-256 `33d85a12...c248`; the full digest is in Section 6). It is stored URI-encoded, 15,560 bytes, and runs in bootloader’s SVG-JS runtime, which provides an SVG root, a seeded generator `rnd()` built from `splitmix64` [Steele et al., 2014] feeding `sfc32` [Doty-Humphrey], and the token’s `iterationNumber`. Line numbers below refer to the stored code.

The grids. Each photograph was converted to grayscale, center-cropped to the grid’s 28×38 aspect (close to the sheet’s 3×4), autocontrasted with a 1% cutoff, resampled to 28×38 cells, and quantized to 16 levels, written one hexadecimal character per cell: 1,064 bytes per photograph, carrying at most 532 bytes of information. Both strings are literals in the code (lines 25 and 26). Storing the two cost $2,128 \times 250$ mutetz = 0.532 tez. Steyerl’s poor image gives up resolution to circulate [Steyerl, 2009]; these grids gave it up to fit in contract storage. The grids correlate with each other at $r = 0.10$ over cells. Figure 2 shows both.

The memory field. Let $P(x, y) \in [0, 1]$ be the bilinear interpolation of the token’s grid at canvas position (x, y) , where 1 is the lightest level, and let $n(x, y)$ be two octaves of value noise [Perlin,

trait	value	weight in code	observed (20,000 seeds)
density	air / field / full	50 / 38 / 12%	49.7 / 38.1 / 12.3%
inks	1 / 2 / 3 / 4	4 / 56 / 30 / 10%	3.8 / 56.2 / 29.9 / 10.1%
source	blue tree / natural gradient	87 / 13%	86.8 / 13.2%
paper	night (carbon)	8%	8.1%
ember	present	55%	54.8%
memory	edition-determined	Eq. (2)	n/a

Table 1: Seeded traits. Observed shares from `repro/census.mjs` at edition 1; the largest deviation is 1.44 standard errors (one ink). The five non-night papers (blush, ivory, rose, fog, overhead blue at 28/22/20/12/10%) are not reported in the trait line and were not counted.

1985] over a 16×16 lattice drawn from the first 256 outputs of `rnd()` (lines 35 to 42). The generator draws from a darkness field

$$D(x, y) = f(1 - P(x, y)) + 0.85(1 - f)n(3.1\frac{x}{W} + 5.2, 3.1\frac{y}{H}), \quad (1)$$

with memory factor

$$f(e) = \max\left(0.12, 1 - 0.88 \frac{\min(e, 44) - 1}{43}\right) \quad (2)$$

for edition number e (line 69). Over $e = 1, \dots, 44$ the floor never binds: f falls by $0.88/43 \approx 0.0205$ per edition from 1 to exactly 0.12. Platform previews use $f = 1$.

Columns and dots. The canvas is divided into columns of width $3r_0$, where r_0 is a base dot radius drawn per token; the factor 3 and the column-of-dots loop are Haskell’s [Haskell, 2022]. Whether a column is drawn, and how far down it runs, depends on a noise value per column and on D at the column’s center. Down each column the generator places dots of radius $r = r_0(0.72 + 0.56n')(0.8 + 0.5D)$, so dark regions of the memory field get larger dots and longer runs. Each dot’s ink is chosen from D at its center: the token’s k inks are sorted from light to dark and indexed by $\lfloor kD \rfloor$; below $D = 0.1$ the dot is usually bone paper color, and above 0.88 usually carbon. With probability 0.58 a dot repeats the previous dot’s ink when D has changed by less than 0.3 since that dot (line 129), which turns per-dot color noise into runs that break where the picture changes. On night paper the ink order is reversed. A static analogue is the column-structured ordered dither of print halftoning [Ulichney, 1987], with the dither threshold replaced by a field whose photographic weight falls with the edition number.

Inks and traits. The five inks are colors I had already used in earlier on-chain works and listed in a palette file before this project began: pink and pale blue from “together”, two blues from “overhead”, and a pale blue labeled “mr. blue sky” whose source work I can no longer identify. Bone and carbon come from earlier noise studies, and the single ember dot (`#e0521b`) from the fourth 2022 test output. Table 1 lists the seeded traits with the weights written in the code and their frequencies over 20,000 seeds; every observed share is within 1.5 standard errors of its weight. Memory is the only trait not drawn from the seed. Each token writes its traits, including memory as a percentage, into the SVG’s `desc` element (line 153).

4 What the editions remember

The trait line asserts a memory percentage, which is $f(e)$ and nothing more. Whether a viewer, or a program, could recover the photograph from a finished edition is a different question, and the answer depends on everything between Eq. (1) and the drawn dots. I measured it at two stages.

Method. For the field, I rebuilt n for each seed from the first 256 draws of its PRNG, evaluated Eq. (1) on a 75×100 lattice, and computed the Pearson correlation between D and the photograph’s darkness $1 - P$, for both photographs at every edition (`repro/field.mjs`). Neither measure is an information bound. The photograph is never lost from the work: both grids sit in the code, and with the seed’s noise any edition’s field can be inverted exactly, $1 - P = (D - 0.85(1 - f)n)/f$. What the measures track is how strongly the photograph shows in the field and in the drawing. For the drawing, I rendered each (seed, edition) pair with the stored generator, took every dot except the ember, and computed the Spearman rank correlation between the dot’s ink darkness (one minus Rec. 601 luma) and the photograph’s darkness at the dot’s center (`repro/measure.mjs`). Centers are taken before the texture filter, which displaces each dot by at most 3 pixels, about a ninth of a grid cell. A token’s ink darkness takes at most $k + 2$ values (k inks, bone and carbon), so ties alone keep the rank correlation below 1 even at full memory. Night-paper tokens invert the ink order by design, so their correlations are sign-flipped; at edition 1 their median is 0.80, against 0.83 over all 1,000 seeds. The measure reads ink only, although dot size, column presence and run length also follow D ; a second measure below counts those too. In 3,291 renders (7.5%) every dot has the same ink and the rank correlation is undefined. I score these 0, since their ink carries no rank information about the photograph, so every render counts. For tree tokens their share rises from 1.0% at edition 1 to 13.3% at edition 42 (12.7% at edition 44); dropping them instead would raise the edition-44 median from 0.14 to 0.19. Seeds are $1,000,003 + 7,919i$ for $i = 1, \dots, 1000$, giving 44,000 renders, 872 of each edition from tree tokens and 128 from sky tokens. Because the PRNG stream depends on the seed alone, the 44 renders of one seed share every trait drawn before the dots (all but the ember); the curves pool over seeds, not over editions.

A counterfactual null. For each (seed, edition) I also render the same program with only the photograph swapped: the source draw still consumes its random number, and every trait except the ember is drawn before the dots, so those traits come out the same; in 600 checked pairs they all matched (the ember is drawn after the dots and can differ) (`renderSwapped` in `repro/runtime.mjs`; the stored code is not modified). Inside the dot loop, dot sizes, run lengths and the random draws follow the other photograph. The dots of that render are correlated with the photograph the real token sampled. This is a replacement control: it shows what the same token would look like if the other photograph drove it, not what a token with no photograph would look like.

A closed form for the field. If the noise is uncorrelated with the photograph, Eq. (1) gives

$$\rho_{\text{field}}(f) \approx \frac{f \sigma_p}{\sqrt{f^2 \sigma_p^2 + 0.85^2 (1 - f)^2 \sigma_n^2}}, \quad (3)$$

where σ_p is the standard deviation of the photograph’s darkness over the canvas and σ_n that of the noise. Measured (`repro/sigma.mjs`): $\sigma_p = 0.2515$ for the tree and 0.2601 for the sky; the median σ_n over 1,000 seeds is 0.1486. The photograph and the noise contribute equal spread where $f \sigma_p = 0.85(1 - f) \sigma_n$, at $f^* = 0.334$ for the tree and 0.327 for the sky, which Eq. (2) reaches at edition




Figure 3: How much of the photograph survives, by edition. The field curve is a Pearson correlation over the whole canvas; the dot curves are Spearman correlations over drawn dots, so the two share a shape but not a scale. Eq. (3) is omitted because it lies within 0.010 of the tree’s field curve. Medians over 1,000 seeds, every render counted (872 tree and 128 sky renders per edition); the band is the tree’s interquartile range.

33.5 and 33.9. At f^* the correlation is $1/\sqrt{2} \approx 0.71$. At edition 34 ($f = 0.325$) Eq. (3) gives 0.69 and 0.70 and the measured medians are 0.70 and 0.72; at edition 44 it gives 0.26 and 0.27 against 0.27 and 0.28. Over all 44 editions the closed form is within 0.010 of the tree’s measured curve (largest gap at edition 32) and 0.021 of the sky’s (largest gap 0.0204, at edition 38), and the measurement sits at or above it at every edition but one. The gap comes from the covariance between the noise and the photograph, which is zero only on average over seeds and enters both the numerator and the denominator of the exact correlation, through $\text{cov}(D, 1 - P) = f\sigma_p^2 + 0.85(1 - f)\text{cov}(n, 1 - P)$. Recomputing the medians with each seed’s own σ_n and the covariance set to zero reproduces Eq. (3) to four decimal places, as expected when a median passes through a monotone function (up to the averaging of the two middle values in a sample of 1,000); restoring each seed’s covariance reproduces the whole gap (`review/probes-r2.mjs`).

Results. Figure 3 and Table 2 give the curves.

Between editions 1 and 22 the field’s correlation with the tree falls by 0.06; between editions 34 and 44 it falls by 0.43, out of a total loss of 0.73. The drawn dots follow the same shape less sharply: a fall of 0.13 between editions 1 and 22 and 0.30 between 34 and 44. Equation (3) gives the reason: while $f\sigma_p$ is large next to $0.85(1 - f)\sigma_n$ the noise barely moves the ranking of dark and light, and the photograph loses its hold only once the two spreads are comparable, around edition 34. Any mixture whose signal weight falls linearly while its noise weight rises holds and then drops; where the drop comes depends on the ratio of the two spreads, here around edition 34. The shape is a property of the mixture, not a discovery about memory. Plotted against edition number it has the opposite outline to human retention curves, which fall fastest first [Ebbinghaus, 1885, Murre and Dros, 2015, Wixted and Ebbesen, 1991], but the axes are not commensurable:

edition	f	field, tree	dots, tree	(IQR)	single-ink %, tree	dots, sky	null, tree
1	1.000	1.000	0.824	(0.741, 0.858)	1.0	0.845	0.067
9	0.836	0.996	0.793	(0.690, 0.835)	2.5	0.814	0.075
15	0.713	0.981	0.753	(0.636, 0.807)	5.5	0.770	0.101
22	0.570	0.938	0.693	(0.545, 0.764)	6.3	0.704	0.060
30	0.407	0.816	0.540	(0.349, 0.659)	8.8	0.528	0.014
34	0.325	0.699	0.433	(0.205, 0.594)	11.0	0.409	0.009
44	0.120	0.267	0.136	(0.000, 0.338)	12.7	0.128	0.000

Table 2: Selected rows of Figure 3. Field: median Pearson correlation of D with the photograph. Dots: median Spearman correlation of drawn ink darkness with the photograph at each dot, over all renders, with single-ink renders (share given) scored 0. Null: the counterfactual render in which the other photograph drives the field, correlated with the tree. Full tables in `data/curve.csv` and `data/field.csv`.

those curves follow one rememberer through time, while each edition’s level is fixed at mint and no token changes afterwards.

At edition 1 the field correlation is exactly 1, and the drawn dots reach a median of 0.82 (tree) and 0.84 (sky). Part of that gap is the tie ceiling described above and part is the drawing: sparse columns, run persistence, and dots placed where D is dark. The two measures use different statistics on different supports (Pearson over the whole canvas, Spearman over the dots), so their levels are not directly comparable; only their shapes are. The drawn measure depends on the traits: the median at edition 1 is 0.77 for air density, 0.84 for field and 0.87 for full, 0.42 for one-ink tokens, which have only bone, carbon and the single ink to choose from, and 0.87 for four-ink tokens. The quantizer also loses levels as the run goes on: carbon needs $D > 0.88$, and from edition 27 no carbon dot appears on light paper in any of the 1,000 seeds, so by the time the steep part of the decline begins, the drawing has one ink level fewer to express it with.

At edition 44 the median drawn correlation is 0.14 for the tree; 62.5% of the 872 tree renders are positive, and 629 exceed their own counterfactual null ($z = 13$ against an even split). The decline is orderly in the medians and not token by token: a random edition-44 tree render scores above a random edition-34 render 26% of the time, and adjacent editions are close to an even split (edition 23 above edition 22: 48%). In a set of 44 tokens actually minted, a later token can easily show more of the photograph than an earlier one. The 12% floor in Eq. (2) leaves a trace a statistic can find, though I would not claim a viewer could.

The counterfactual null is near zero but not exactly: its median by edition ranges from 0.000 to 0.101 for tree tokens and to 0.104 for sky tokens, consistent with the two grids’ shared large-scale structure (cell correlation 0.10). From edition 37 the tree null’s median is exactly 0, because the single-ink null renders, scored 0, fall across the middle of the distribution; the paired comparison above is the better test at those editions. An earlier version of this analysis compared each token’s own dots with the other photograph; because those dots were placed by the sampled photograph, that null ran higher, at 0.05 to 0.13 for tree tokens.

Ink, size and presence. A second measure counts every channel: for each of the 1,064 grid cells, the sum over dots centered in it of dot area times the dot’s darkness above the paper, correlated (Pearson) with the grid, with the same counterfactual null (`repro/extra.mjs`). It is much weaker: 0.20 for the tree at edition 1 against 0.11 for its null, and 0.14 against 0.13 at edition 44, because

most cells hold no dot. At the five editions measured, the median area score of single-ink renders is below the null’s median for tree tokens and at most 0.015 above it for sky tokens (9 to 19 renders each), so scoring their ink measure 0 loses little. The ink measure is the more sensitive of the two, and it is the one reported above.

The label against the measurement. Each token’s trait line prints f , rounded to a whole percent, as its memory. Read against the median drawn measure, the printed label is too high through edition 11 (84% at edition 9 against a tree median of 0.79) and too low from edition 12, for both photographs; the one exception is the sky at edition 42, where the median (0.156) falls just below the label (16%). At edition 44 the label says 12%, against 0.14 for the tree’s dots and 0.27 for its field. These are medians: an individual render can sit on either side of its label at any edition. A weight and a correlation are different quantities, but the token’s own label invites the comparison.

The collection so far. Nine of the 44 editions are minted as of 12 October 2026. The most recent is edition 9, with $f = 0.836$. At that edition number the field’s median correlation is 0.996, and the median drawn correlation over renders at that edition number is 0.79 for tree tokens and 0.81 for sky tokens. These are medians over seeds, not measurements of the minted token, whose rendered seed I could not reproduce (Section 5). By any measure here, the collection has not yet forgotten much. The steep part of the curve lies in editions not yet minted, and whether the piece completes its forgetting depends on whether anyone collects them.

5 What the ledger remembers

The work assigns permanence to the ledger and loss to the editions. Two months of running it tested both halves of that assignment.

Superseded versions. The generator’s storage key has been written 16 times: once to create it, five times to publish versions 2 to 6, once to open the sale, and nine times to record a mint. All 16 writes are listed in `data/generator-history.csv`. Version 1, published at 22:10:28 UTC on 31 July, opened with the comment “in november 2022 i wrote a p5.js sketch called ‘365 days of tbm’ ... this is that algorithm, rewritten to run entirely on-chain.” Edition 1 was minted under it four minutes later. That evening I asked for a token-level comparison of my 2022 sketch with Haskell’s tutorial: the function that draws the columns was 95% identical to his by token count, and two helpers were identical character for character. Version 4 (03:03 UTC, 1 August) rewrote the header to credit him, versions 5 and 6 added a corrected description, and the public working log says what happened [Meadows, 2026c]. The contract can overwrite the current value of the key, which is what a version update does, but every earlier value remains in the block that wrote it and the TzKT indexer still returns all six versions [Baking Bad, 2026]; `repro/verify-chain.mjs` retrieves them. Version 1 survives because archive nodes and indexers keep history, and the false claim stays on that record next to its correction. The mechanism is the one Finck identifies as the obstacle to erasure on ledgers [Finck, 2018], met here from the author’s side.

Three other claims were caught before mint and never reached the chain: that the fifth ink came from the sky photograph (the palette file says otherwise), and twice that neither photograph had been published (both were minted on Ethereum in 2023; the sky file I believed unpublished

matched “Natural Gradient” in capture time, GPS position and exposure). The rule I took from them is that the work claims only what a public record can confirm.

Seeds. The ledger keeps the program and, for each token, an edition number and the entropy the platform drew at mint. That should make every edition reproducible from public data, and with the documented procedure it does not. For edition 1, the seed the token page displays, the seed the platform’s viewer renders from, and the seed declared in the stored artifact are three different values; the viewer reduces the seed modulo 2^{32} before rendering, and the same mismatch appears on other artists’ tokens. I reported this to the platform’s maintainers [Meadows, 2026a], and it is unresolved as of this writing. Separately, a regeneration call wrote raw seed bytes into edition 1’s artifact URI, which is therefore not valid UTF-8. The consequence for this work is specific. The memory factor of every token is recoverable from public data alone, because it depends only on the edition number. The platform’s published artifact template seeds its generator exactly as `repro/runtime.mjs` does (splitmix64 into sfc32). Rendering edition 1 from its stored entropy, from that entropy reduced modulo 2^{32} , and from a viewer seed I recorded in August gives three different pictures, none of them the platform’s own image of the token. Either the platform renders from a seed it does not store on chain, or its image predates a regeneration of the token; public data does not say which. The only image of edition 1 I can vouch for is the platform’s cached render. I have checked edition 1 in detail, not every token.

Copies off the ledger. The work’s project page fetches the generator from the contract and compares it with the copy it runs, printing “Verified” or “Mismatch” to the reader. In August 2026 a site-wide punctuation pass rewrote six dashes inside comments of the page’s copy, 12 bytes, and the page told every visitor for eight days that its code did not match the chain. That copy is now checked against a pinned digest at every build. This paper’s copy (`repro/essentialism.chain.js`) is checked the same way.

6 Reproducibility

Everything in Sections 3 and 4 is regenerated by `make` in the paper’s directory. The package is archived on Zenodo as <https://doi.org/10.5281/zenodo.23333266> and maintained at <https://github.com/TheBenMeadows/essentialism-preprint>. The pieces:

- `repro/essentialism.chain.js`: the generator, byte-identical to version 6 on Tezos (contract `KT1CB4MYiAViCuXWBU961x7LjQXGeA8SnQwt`, big map 771404, key 407), SHA-256 `33d85a12cf2a8404bccda1e31d6b3fd72c361673f995cc58b0857eabb6d5c248`. `repro/verify-chain.mjs` fetches the stored code through TzKT, checks the digest, and writes the version history; it is the only step that needs the network.
- `repro/runtime.mjs`: a stand-in for the bootloader runtime (splitmix64 into sfc32, and a minimal DOM that serializes SVG). For seed 1,007,922 at edition 1 its output is byte-identical to the renderer used during development, and its PRNG matches the platform’s published artifact template line for line. I could not compare it with the platform’s own renders, because no seed I could derive from chain data reproduces them (Section 5). Results over many seeds do not depend on that comparison; statements about a specific minted token do.

- `repro/census.mjs`, `field.mjs`, `sigma.mjs`, `measure.mjs`, `extra.mjs`, `plotdata.mjs`, `stats.mjs`, `figures.mjs`: Table 1, the field curve, the variances, the dot correlations and counterfactual null, the area-weighted measure and the carbon count, the plot table, the measured numbers quoted in the text (`data/stats.txt`), and Figure 1. All are deterministic; seeds are fixed in the scripts. The full measurement takes under a minute with Node.js 26.11 on an Apple-silicon Mac.
- `data/`: every CSV the paper reads, the raw TzKT responses, and SHA256SUMS.

Three things are not reproducible from this package. I could not rebuild the rendered image of a minted token from public chain data (Section 5); its memory factor can be computed from its edition number. The photographs’ full-resolution files are not included; the published images are linked in the references, and `repro/embed-photo.py` is the script that produced the grids, but the published images may not match the exact crop and export I used, so reapplying it will not reproduce the stored strings byte for byte. The stored strings are the authority. And the figures in Figure 1 depend on Inkscape’s SVG filter implementation for texture; browsers render the same geometry with slightly different grain.

7 Limits

The measurement is statistical. A rank correlation between ink darkness and a grid says that information survives; it does not say that a person would recognize a tree, and I ran no viewer study. The field correlation assumes a uniform lattice over the whole canvas, while the drawn measure only sees places where dots fall, so the two are not directly comparable in level, only in shape. The counterfactual null is not zero, so not every part of a reported correlation is due to the photograph the token sampled; the null does not say how much is. A perceptual measure such as SSIM [Wang et al., 2004] on a rasterized render would answer a nearer question to “could someone see the tree”; I used rank correlation because it needs no rasterizer and no choice of scale. The seeds are an arithmetic sequence rather than a random sample of the 64-bit seed space; the trait frequencies in Table 1 match the coded weights, which is consistent with an unbiased sequence but does not prove one. The 44,000 renders are 1,000 seed trajectories through the 44 editions, not independent draws, and the sky’s drawn-dot curves rest on 128 of them (the field curves use all 1,000 for both photographs). Finally, 35 of the 44 editions do not exist yet. Every statement here about editions 10 to 44 is about what the generator will produce for them, not about tokens anyone holds, and the part of the curve that carries most of the forgetting is entirely in that range.

Acknowledgements

The column structure of this work comes from Andy Haskell’s tutorial [Haskell, 2022], credited in the on-chain code since generator version 4.

Use of AI tools

The port of my 2022 sketch to bootloader’s runtime (the runtime translation, noise function, grid embedding, fade arithmetic and texture filters), the measurement and audit scripts, and the text

of this paper were drafted with Claude (Anthropic) running in Claude Code, directed by me. The photographs, the ink library, the premise of a fading memory, the edition count and the corrections in Section 5 are my decisions, and I take responsibility for all content. Drafts were reviewed by GPT-6 Astra, Claude Fable 5.1, SWE-2, Kimi K3 and Grok 4.6 through Cognition’s Devin CLI, by DeepSeek V4.1 Flash and GPT-6 Luna, and by Gemini; every measurable defect a reviewer reported was reproduced or refuted by a script before it changed the paper, and the probe and audit scripts that check the measured numbers and the chain facts in the text against the data are in the repository. Every reference was checked against the work itself or a published description of it.

References

- Baking Bad. TzKT: Tezos blockchain indexer and API. <https://api.tzkt.io>, 2026. Accessed 2026-10-12.
- Liam J. Bannon. Forgetting as a feature, not a bug: The duality of memory and implications for ubiquitous computing. *CoDesign*, 2(1):3–15, 2006. doi: 10.1080/15710880600608230.
- Roland Barthes. *Camera Lucida: Reflections on Photography*. Hill and Wang, New York, 1981.
- William Basinski. The Disintegration Loops. 2062 Records, 2002. Audio recordings, four volumes, 2002–2003.
- Juan Benet. IPFS - content addressed, versioned, P2P file system, 2014.
- Jorge Luis Borges. Funes el memorioso. In *Ficciones*. Editorial Sur, Buenos Aires, 1944. First published in *La Nación*, 1942.
- Ian Burn. Xerox Book. Artist’s book: 100 successive photocopies of a blank sheet on a Xerox 720, bound in the order made, 1968.
- Wendy Hui Kyong Chun. The enduring ephemeral, or the future is a memory. *Critical Inquiry*, 35(1):148–171, 2008. doi: 10.1086/595632.
- Court of Justice of the European Union. Google Spain SL and Google Inc. v Agencia Española de Protección de Datos (AEPD) and Mario Costeja González, Case C-131/12. ECLI:EU:C:2014:317, 2014.
- Department of Drawings and Prints. The printed image in the west: Drypoint. Heilbrunn Timeline of Art History, The Metropolitan Museum of Art, <https://www.metmuseum.org/essays/the-printed-image-in-the-west-drypoint>, 2003. October 2003.
- Chris Doty-Humphrey. PractRand: a C++ library of pseudo-random number generators and statistical tests, including the sfc32 generator. <https://pracrand.sourceforge.net/>. Software.
- Hermann Ebbinghaus. *Über das Gedächtnis: Untersuchungen zur experimentellen Psychologie*. Duncker & Humblot, Leipzig, 1885.
- Wolfgang Ernst. *Digital Memory and the Archive*. University of Minnesota Press, Minneapolis, 2013.

European Parliament and Council of the European Union. Regulation (EU) 2016/679 (General Data Protection Regulation), Article 17: Right to Erasure. Official Journal of the European Union, L 119, 4 May 2016, 2016.

Michèle Finck. Blockchains and data protection in the European Union. *European Data Protection Law Review*, 4(1):17–35, 2018. doi: 10.21552/edpl/2018/1/6.

Sarah Friend. Lifeforms. Tokens on Polygon, contract 0x8916edd9b39783d85303ecc6613917ddd735d88d. <https://lifeforms.supply/>, 2021.

John Gerrard. World Flag. Art Blocks and Pace, contract 0xea698596b6009a622c3ed00dd5a8b5d1cae4fc36, project 6, 2023.

L. M. Goodman. Tezos: A self-amending crypto-ledger. <https://tezos.com/whitepaper.pdf>, 2014. White paper.

Andy Haskell. Convert images to mosaics in p5.js. <https://dev.to/andyhaskell/convert-images-to-mosaics-in-p5js-2dlc>, January 2022. Tutorial; finished sketch at https://editor.p5js.org/andyhaskell/sketches/fwaWSmAh_.

Matthew G. Kirschenbaum. *Mechanisms: New Media and the Forensic Imagination*. MIT Press, Cambridge, MA, 2008.

Martin Klein, Herbert Van de Sompel, Robert Sanderson, Harihar Shankar, Lyudmila Balakireva, Ke Zhou, and Richard Tobin. Scholarly context not found: One in five articles suffers from reference rot. *PLoS ONE*, 9(12):e115253, 2014. doi: 10.1371/journal.pone.0115253.

Larva Labs. Autoglyphs. Ethereum, contract 0xd4e4078ca3495de5b1d4db434bebc5a986197782. <https://www.larvalabs.com/autoglyphs>, 2019.

Alvin Lucier. I Am Sitting in a Room, 1970. Sound work for voice and electromagnetic tape. Recording: Lovely Music VR 1013 (LP), 1981.

Viktor Mayer-Schönberger. *Delete: The Virtue of Forgetting in the Digital Age*. Princeton University Press, Princeton, NJ, 2009.

Benjamin Meadows. Blue Tree. Photograph, 1 of 1, minted on Ethereum 2023-09-06. Contract 0x93ad4670D8Dd9FC3987EB296315827C3D138ADeA, token 5. <https://eth.blockscout.com/token/0x93ad4670D8Dd9FC3987EB296315827C3D138ADeA/instance/5>, 2023a.

Benjamin Meadows. Natural Gradient. Photograph, 1 of 1, minted on Ethereum 2023-09-17, metadata on Arweave. Contract 0x93ad4670D8Dd9FC3987EB296315827C3D138ADeA, token 6. <https://eth.blockscout.com/token/0x93ad4670D8Dd9FC3987EB296315827C3D138ADeA/instance/6>, 2023b.

Benjamin Meadows. Seed handling: regenerate_token stores raw bytes in artifactUri (invalid JS/UTF-8); stored artifact seed differs from frontend render seed. <https://github.com/objkt-com/bootloader-monorepo/issues/30>, 2026a. GitHub issue, opened 2026-07-31.

Benjamin Meadows. essentialism. Generative work, bootloader svg-js generator 407, contract KT1CB4MYiAViCuXWBU961x7LjQXGeA8SnQwt on Tezos. <https://bootloader.art/generator/svg-js/407>, 2026b. Published 2026-07-31; generator version 6 since 2026-08-01.

- Benjamin Meadows. essentialism: the working log. <https://thebenmeadows.com/blog/essentialism-the-working-log/>, 2026c. Published 2026-07-31, corrected 2026-08-26.
- Gustav Metzger. Auto-Destructive Art. Manifesto, London, 4 November 1959, 1959.
- Jaap M. J. Murre and Joeri Dros. Replication and analysis of Ebbinghaus’ forgetting curve. *PLoS ONE*, 10(7):e0120644, 2015. doi: 10.1371/journal.pone.0120644.
- Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf>, 2008.
- objkt.com. bootloader: on-chain generative art platform on Tezos. [https://bootloader.art; source: https://github.com/objkt-com/bootloader-monorepo](https://bootloader.art;source=https://github.com/objkt-com/bootloader-monorepo), 2026. Accessed 2026-10-12.
- Octez developers. History modes. Octez and Protocol documentation, https://octez.tezos.com/docs/user/history_modes.html, 2026. Accessed 2026-10-12.
- Roman Opalka. OPALKA 1965/1 – ∞ . Series of numbered paintings, 1965–2011. Description of the progressive whitening of the grounds: Christie’s lot essay, <https://www.christies.com/en/lot/lot-6100726>, 1965.
- Ken Perlin. An image synthesizer. *ACM SIGGRAPH Computer Graphics*, 19(3):287–296, 1985. doi: 10.1145/325165.325247.
- Casey Reas. 923 EMPTY ROOMS. Art Blocks and Bright Moments, contract 0x145789247973c5d612bf121e9e4eef84b63eb707, project 1. <https://www.artblocks.io/collection/923-empty-rooms-by-casey-reas>, 2023.
- Richard Rinehart and Jon Ippolito. *Re-collection: Art, New Media, and Social Memory*. Leonardo. MIT Press, Cambridge, MA, 2014. doi: 10.7551/mitpress/7975.001.0001.
- David S. H. Rosenthal. Keeping bits safe: How hard can it be? *Communications of the ACM*, 53(11):47–55, 2010. doi: 10.1145/1839676.1839692.
- John F. Simon, Jr. Every Icon. Java applet, unlimited numbered edition. <https://artport.whitney.org/v1/everyicon.html>, 1997.
- Guy L. Steele, Doug Lea, and Christine H. Flood. Fast splittable pseudorandom number generators. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications (OOPSLA)*, pages 453–472, 2014. doi: 10.1145/2660193.2660195.
- Hito Steyerl. In defense of the poor image. *e-flux journal*, (10), 2009. <https://www.e-flux.com/journal/10/61362/in-defense-of-the-poor-image/>.
- Robert Ulichney. *Digital Halftoning*. MIT Press, Cambridge, MA, 1987.
- Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. doi: 10.1109/TIP.2003.819861.
- Sam Williams, Viktor Diordiiev, Lev Berman, India Raybould, and Ivan Uemlianin. Arweave: A protocol for economically sustainable information permanence. <https://www.arweave.org/yellow-paper.pdf>, 2019. Yellow paper, draft 1.

- John T. Wixted and Ebbe B. Ebbesen. On the form of forgetting. *Psychological Science*, 2(6): 409–415, 1991. doi: 10.1111/j.1467-9280.1991.tb00175.x.
- Gavin Wood. Ethereum: A secure decentralised generalised transaction ledger. <https://ethereum.github.io/yellowpaper/paper.pdf>, 2014. Yellow paper, first published 2014; the URL serves the current revision.
- Jonathan Zittrain, Kendra Albert, and Lawrence Lessig. Perma: Scoping and addressing the problem of link and reference rot in legal citations. *Harvard Law Review Forum*, 127:176–199, 2014.