Backtrace: The Mathematics of an Honest Wildfire-Origin Estimator

Charlie Ramus

July 21, 2026

Abstract

Backtrace is an offline field instrument for reasoning about where a wildland fire started. An investigator walks the burn, flags physical fire-pattern indicators, and records the direction each one points; Backtrace fuses those bearings — each with its own angular uncertainty — into a probability field over the origin. Because peer-reviewed testing finds that fire-pattern indicators carry roughly **103°** of mean directional error, the estimator is designed to *refuse false confidence*: circular von Mises likelihoods evaluated on a planar grid posterior, a uniform outlier mixture, Bayesian macro priors, and highest-density credible *regions* rather than a point. This note explains the project and derives the mathematics behind it.

1 The problem, and why it forbids a naive solution

Wildfire-origin doctrine (NWCG PMS 412) reads physical artifacts — the angle of char on a stem, soot deposited on the fire-facing side of a rock, an unburned "protection" shadow — as arrows pointing back toward the fire's origin. Each indicator i, standing at position $\mathbf{p}_i$, yields an observed back-azimuth θ_i. The obvious estimator is to draw the rays and intersect them.

The validation literature says this cannot work. Parker & Babrauskas (2024) had four court-qualified senior investigators assess 404 indicators across three test burns with drone-verified ground truth. The mean absolute directional error was **103°** — against a random-guessing baseline of 90°. An error band that wide will not triangulate. Any tool that draws intersecting rays and prints the crossing point is manufacturing confidence that the data does not contain.

Backtrace's answer is to model that error honestly. Each indicator type carries an empirical angular standard deviation σ_i taken from the published per-indicator results, which is converted into a concentration parameter κ_i of a circular distribution (Section 3). Table 1 shows the priors the app ships with.

Indicator	n	$\hat{\sigma}$ **(deg)**	$R = e^{-\sigma^2/2}$	κ
Protection	39	81	0.368	0.79
White ash	6	81	0.368	0.79
Sooting	20	97	0.239	0.49
Angle of char	89	98	0.232	0.48
Grass stem	7	98	0.232	0.48
Staining	133	106	0.181	0.37

Table 1: Per-indicator directional error (Parker & Babrauskas 2024, Table 5) and the von Mises concentration κ it implies. A κ below one is a nearly flat distribution: the math itself refuses to be confident.

2 Geometry: all math on a local tangent plane

An origin search area is 10 m to 1 km across. Over 1 km the difference between a great-circle bearing and a plane bearing is under $0.01°$, so spherical trigonometry buys nothing but bugs. Backtrace instead anchors a local East-North-Up (ENU) tangent plane at a point A and converts every WGS84 coordinate into metres on that plane, exactly: geodetic $\to$ ECEF $\to$ ENU. With semi-major axis $a = 6378137$ m, flattening $f = 1/298.257223563$ and $e^2 = f(2-f)$, the prime-vertical radius and ECEF coordinates are

$$N(\varphi) = \frac{a}{\sqrt{1 - e^2 \sin^2 \varphi}}, \qquad \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} (N+h)\cos\varphi\cos\lambda \\ (N+h)\cos\varphi\sin\lambda \\ \big(N(1-e^2)+h\big)\sin\varphi \end{pmatrix}$$

and the ENU coordinates of a point $\mathbf{p}$ relative to the anchor $\mathbf{p}_0$ are the rotated difference

$$\begin{pmatrix} E \\ N \\ U \end{pmatrix} = \begin{pmatrix} -\sin\lambda_0 & \cos\lambda_0 & 0 \\ -\sin\varphi_0\cos\lambda_0 & -\sin\varphi_0\sin\lambda_0 & \cos\varphi_0 \\ \cos\varphi_0\cos\lambda_0 & \cos\varphi_0\sin\lambda_0 & \sin\varphi_0 \end{pmatrix} (\mathbf{p} - \mathbf{p}_0).$$

Azimuths are measured clockwise from true north, so a bearing α becomes the unit vector $\mathbf{d} = (\sin\alpha,\ \cos\alpha)$ in (E, N) — note the sine on East. All geometry below happens in these planar metres; latitude and longitude reappear only at the display and export boundary.

3 The likelihood: von Mises on a grid

Angular error is circular — $359°$ and $1°$ are two degrees apart, not 358 — so the noise model is the von Mises distribution, the circular analogue of the Gaussian. Given the PDF of a von Mises distribution:

$$p(\theta \mid \mu, \kappa) \;=\; \frac{e^{\kappa \cos(\theta-\mu)}}{2\pi\, I_0(\kappa)}$$

where I_0 is the modified Bessel function of order zero and $\kappa \ge 0$ is the concentration ($\kappa = 0$ is uniform on the circle; large κ approaches a Gaussian of variance $1/\kappa$), we can score every candidate origin. Rasterise the search region into cells $\mathbf{x}$ (a 1 km area at 2 m cells is $500 \times 500 = 250{,}000$ cells). For each cell and each indicator node i, the *expected* back-azimuth if the fire had truly started at $\mathbf{x}$ is

$$\beta_i(\mathbf{x}) = \operatorname{atan2}\!\big(x_E - p_{i,E},\ x_N - p_{i,N}\big), \qquad \delta_i = \operatorname{wrap}_{\pm\pi}\!\big(\theta_i - \beta_i(\mathbf{x})\big),$$

and the likelihood that node i observed θ_i given origin $\mathbf{x}$ is $p(\theta_i \mid \mathbf{x}) = e^{\kappa_i \cos\delta_i} / \big(2\pi I_0(\kappa_i)\big)$. Because $\cos\delta_i$ is negative for anything more than 90° off, a cell directly behind the observer is penalised by $e^{-\kappa_i}$ automatically — the ray constraint falls out of the model for free.

3.1 From a measured σ to κ

Each node stores a circular standard deviation σ_i (radians) — from the published indicator priors of Table 1, from the live capture window's circular statistics, or from the two-point GNSS bearing's propagated accuracy. The mean resultant length is $R = e^{-\sigma^2/2}$, and κ solves $R = I_1(\kappa)/I_0(\kappa)$, inverted with Fisher's (1993) approximation:

$$\kappa \;\approx\; \begin{cases} 2R + R^3 + \frac{5}{6}R^5 & R < 0.53 \\ -0.4 + 1.39\,R + \dfrac{0.43}{1-R} & 0.53 \le R < 0.85 \\ \dfrac{1}{R^3 - 4R^2 + 3R} & R \ge 0.85 \end{cases}$$

Sanity check: $\sigma = 90^\circ = 1.571$ rad gives $R = 0.29$ and $\kappa \approx 0.61$ — nearly flat. A node this uncertain *cannot* shrink the candidate region, and in this formulation it provably does not.

3.2 Robustness and GPS uncertainty

Roughly 15% of indicators are grossly misread, so each likelihood is mixed with a uniform outlier component,

$$L_i(\mathbf{x}) \;=\; (1-\varepsilon)\,\frac{e^{\kappa_i \cos\delta_i}}{2\pi I_0(\kappa_i)} \;+\; \frac{\varepsilon}{2\pi}, \qquad \varepsilon = 0.15,$$

which bounds the influence of any single wildly-wrong node — a one-line, principled replacement for RANSAC. Position error inflates the angular error by the subtense of the GPS accuracy at working range r_i: $\sigma_{i,\text{eff}} = \sqrt{\sigma_i^2 + (h_{\text{acc}}/r_i)^2}$ (at 5 m accuracy and 30 m range, that adds 9.5°).

4 Fusion: the Bayesian grid posterior

Macro evidence — a V-pattern apex, the burn perimeter, a witness first-smoke cone, a dispatch first-report location, an exclusion zone — is region-shaped, not ray-shaped, so it enters as a *prior* $\pi_0(\mathbf{x})$ rather than a likelihood. The posterior over the grid is the doctrine-shaped Bayes update, computed in log space:

$$\log \pi(\mathbf{x} \mid \theta_{1:n}) \;=\; \log \pi_0(\mathbf{x}) \;+\; \sum_{i=1}^{n} \log L_i(\mathbf{x}) \;+\; C,$$

normalised so the grid sums to one. With no macro constraints the prior is uniform and the result is identical to the likelihood-only estimate — a hard invariant in the code. The update is incremental for free: adding a node is one multiply into the existing posterior, which is the recursive property people want from a Kalman filter, without the Kalman assumptions (a Kalman filter estimates a *time-evolving* state under *additive Gaussian* noise; the origin is static and the noise is circular and heavy-tailed). Cost: $250{,}000$ cells $\times$ 40 nodes = 10 M cosine evaluations — about a second in plain JavaScript, made interactive with a coarse-to-fine pyramid.

Table 2 shows the structure of the discretised posterior for a small grid: every cell holds the fused probability that the origin lies within it, exactly as a joint distribution table tabulates outcomes.

	x_E	**0**	**1**	**2**	**3**	**4**	**5**
x_N	$\pi(\mathbf{x}_{j,k})$	$k=0$	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$
0	$j=0$	$\pi_{0,0}$	$\pi_{0,1}$	$\pi_{0,2}$	$\pi_{0,3}$	$\pi_{0,4}$	$\pi_{0,5}$
1	$j=1$	$\pi_{1,0}$	$\pi_{1,1}$	$\pi_{1,2}$	$\pi_{1,3}$	$\pi_{1,4}$	$\pi_{1,5}$
2	$j=2$	$\pi_{2,0}$	$\pi_{2,1}$	$\pi_{2,2}$	$\pi_{2,3}$	$\pi_{2,4}$	$\pi_{2,5}$
3	$j=3$	$\pi_{3,0}$	$\pi_{3,1}$	$\pi_{3,2}$	$\pi_{3,3}$	$\pi_{3,4}$	$\pi_{3,5}$
4	$j=4$	$\pi_{4,0}$	$\pi_{4,1}$	$\pi_{4,2}$	$\pi_{4,3}$	$\pi_{4,4}$	$\pi_{4,5}$
5	$j=5$	$\pi_{5,0}$	$\pi_{5,1}$	$\pi_{5,2}$	$\pi_{5,3}$	$\pi_{5,4}$	$\pi_{5,5}$

Table 2: The discretised grid posterior. Cell (j,k) holds $\pi_{j,k} = \pi(\mathbf{x}_{j,k} \mid \theta_{1:n}) \propto \pi_0(\mathbf{x}_{j,k}) \prod_i L_i(\mathbf{x}_{j,k})$, the fused probability that the origin lies in that cell. The real grid is 500×500 at 2 m resolution.

5 Extracting the answer honestly: highest-density regions

The output is never a point. It is a *highest density region* (HDR): the smallest set of cells containing a stated fraction of the posterior mass. For credibility level c, sort the cell masses in descending order, accumulate until the running sum reaches c, and take the mass of the last cell admitted as the threshold t_c:

$$\mathrm{HDR}_c \;=\; \{\,\mathbf{x} : \pi(\mathbf{x}) \geq t_c\,\}, \qquad \text{where } t_c \text{ satisfies} \sum_{\pi(\mathbf{x})\geq t_c} \pi(\mathbf{x}) = c.$$

Marching squares traces the contour at t_c, and Backtrace renders nested bands at $c = 0.50,\ 0.68,\ 0.95$. Three honesty properties come for free:

Multimodality. If the indicators support two candidate origins, the HDR is simply two polygons. An error ellipse would have averaged them into a lie; the region count n_{modes} is surfaced directly, mapping onto NFPA 921's requirement to consider alternate hypotheses.

Area as a quality number. The 95% region's area in m^2 is reported verbatim. The bundled Marshall Fire demo yields an honestly broad $\sim$19 million m^2 region that *contains* the true origin — precisely because $\kappa \approx 0.5$ refuses to pretend otherwise.

Flatness. The posterior entropy $H = -\sum_{\mathbf{x}} \pi(\mathbf{x}) \log \pi(\mathbf{x})$ is shown as a spread meter: high entropy means "the data says little," and the app says so.

6 Field bearings: where θ_i and σ_i come from

Two-point GNSS bearing (primary, magnetometer-free). Stand at the indicator, take an averaged GPS fix, walk 15–30 m along the indicated direction, take a second fix; the geodesic azimuth between fixes is the bearing. With baseline L and per-fix accuracy h_{acc}, the propagated angular error is approximately

$$\sigma_{2\text{pt}} \;\approx\; \arctan\!\left(\frac{h_{\text{acc}}\sqrt{2}}{L}\right) \;\approx\; 12^\circ \quad (L = 20\ \mathrm{m},\ h_{\text{acc}} = 3\ \mathrm{m}),$$

comparable to a good magnetometer and immune to every magnetic failure mode.

Compass capture (caveated). A $\sim$2 s window of headings is reduced by circular statistics: with $C = \overline{\cos\theta_k}$ and $S = \overline{\sin\theta_k}$, the azimuth is $\operatorname{atan2}(S, C)$, the resultant length is $R = \sqrt{C^2 + S^2}$, and the persisted per-node uncertainty is the circular standard deviation $\sigma_{\text{circ}} = \sqrt{-2\ln R}$ — the very σ_i that feeds κ_i. A disagreement of more than 15° against the two-point bearing flags local magnetic interference.

Declination. Magnetic azimuths convert to true with $\alpha_{\text{true}} = \alpha_{\text{mag}} + D$, where D is computed on-device from the World Magnetic Model 2025 — a degree-12 spherical-harmonic expansion of the geomagnetic potential,

$$V(r, \varphi', \lambda, t) \;=\; a \sum_{n=1}^{12} \sum_{m=0}^{n} \left(\frac{a}{r}\right)^{n+1} \left[\, g_n^m(t) \cos m\lambda + h_n^m(t) \sin m\lambda \,\right] \bar{P}_n^m(\sin\varphi'),$$

with $D = \mathrm{atan2}(Y, X)$ from the resulting field components. Backtrace's implementation matches the NOAA reference calculator to under 0.01°, and stores the raw magnetic azimuth, D, the model epoch, and the derived true azimuth *separately* — a chain-of-custody requirement.

7 Why not least squares

The textbook estimator minimises summed squared perpendicular distances to the bearing lines. With unit normal $\mathbf{n}_i$ to each ray, it has a closed form:

$$\hat{\mathbf{x}} \;=\; M^{-1}\mathbf{b}, \qquad M = \sum_i w_i \, \mathbf{n}_i \mathbf{n}_i^\top, \quad \mathbf{b} = \sum_i w_i \, \mathbf{n}_i \mathbf{n}_i^\top \mathbf{p}_i.$$

It is fast and biased. The noise lives in the *angle*, not the perpendicular offset — a 10° error produces a 3.5 m residual at 20 m but 69 m at 400 m, so equal weighting silently over-trusts distant nodes (Stansfield 1947; Gavish & Weiss 1992). It solves for infinite lines rather than rays, and a single gross outlier drags $\hat{\mathbf{x}}$ arbitrarily far. The grid posterior of Section 4 has none of these defects, at the cost of evaluating cosines on a raster — cheap on modern hardware. And it is a substrate: the planned forward model (v3) replaces the straight-line $\beta_i(\mathbf{x})$ with the direction of ∇T at $\mathbf{p}_i$ from a minimum travel-time solve over a slope-aware Rothermel rate-of-spread raster, and *nothing else in the estimator changes*.

References

Parker, K. & Babrauskas, V. (2024). Validation of NWCG Wildfire Directional Indicators in Test Burns in Coastal California. *Fire* 7(1), 5. DOI 10.3390/fire7010005.

Fisher, N. I. (1993). *Statistical Analysis of Circular Data.* Cambridge University Press. (Von Mises distribution, circular moments, the κ inversion.)

Karney, C. F. F. (2013). Algorithms for geodesics. *J. Geodesy* 87(1), 43–55. (GeographicLib; geodesic azimuths for the two-point bearing.)

NOAA NCEI / BGS (2024). *The World Magnetic Model 2025.* Epoch 2025.0, degree-12 spherical harmonic model; coefficients and reference software.

NWCG (2016). *Guide to Wildland Fire Origin and Cause Determination*, PMS 412. (Indicator taxonomy; GOA → SOA doctrine.)

NFPA 921 (2024 ed.). *Guide for Fire and Explosion Investigations.* (Scientific-method loop; alternate-hypothesis requirement.)

Stansfield, R. G. (1947); Gavish, M. & Weiss, A. J. (1992). *IEEE Trans. AES* 28(3). (Bias of least-squares bearings-only localisation.)

Rothermel, R. C. (1972). *A mathematical model for predicting fire spread in wildland fuels.* USDA INT-115. (Forward model, planned v3.)