The Developmental Morphospace of Wildfire

Ty Tuff, Ph.D.[1]

[1]Environmental Data Science Innovation & Impact Lab (ESIIL), Cooperative Institute for Research in Environmental Sciences (CIRES), University of Colorado Boulder, Boulder, Colorado, USA.

ORCID: 0000-0001-5249-5197

Correspondence: ty.tuff@colorado.edu

One-Sentence Summary

Fire VASE makes the ordered allocation of observed wildfire growth a robust, comparable developmental response.

Abstract

Final area and duration describe where a wildfire ended, but not how it developed. Fire VASE instead represents how observed growth is distributed through relative developmental time. From an archive of 278,569 FIRED events, 10,246 had at least three observations with no intervening calendar-day gaps and defined the primary developmental morphospace. We fitted this space from 20 normalized growth-allocation bins; final area, duration, observation count, absolute growth, and weather were excluded. Broad developmental axes persisted under stricter observation requirements and under Hellinger geometry, although local neighborhoods and extreme exemplars were less stable. Permuting the same increments within fires changed front-loading and pulse summaries, revealing temporal asymmetry relative to that specified null. Among 9,212 weather-complete primary fires, the tested daily-weather summaries and ridge models had weak, response-dependent associations with morphology. For 87,944 exact calendar-day transitions, an autoregressive fire-state baseline provided held-out $R^2 \approx 0.448$; weather added about 0.005, and interactions between weather and state added 0.015–0.018. Fire

VASE therefore provides a developmental representation whose broad geometry is robust under the tested observation-depth, temporal-order, and compositional-geometry comparisons and can be used to investigate why adequately observed fires follow different growth pathways.

1 Introduction

A wildfire's final perimeter is the endpoint of a process. It records how much land burned, but not whether that area accumulated in one rapid run, through weeks of steady expansion, after a long delay, or in repeated episodes of growth and quiescence. Two fires can end at similar sizes after similar durations and still have little else in common. For communities, responders, and ecosystems, those different routes to the endpoint are often the consequential part of the event.

Recent work underscores the importance of growth itself. Across the United States, fast fires—defined as those growing more than 16.2 km^2 in one day—made up 2.7% of analyzed events but accounted for 88% of residential structures destroyed from 2001 to 2020 [1]. Yet the quantities commonly used to compare fires—final area, duration, severity, and mean or peak growth—still reduce the intervening sequence to one or a few numbers. They tell us that growth occurred, but not how it was organized through time.

Weather is an obvious candidate for organizing that sequence. Large-fire activity in western U.S. forests increased with warmer springs and summers and earlier spring snowmelt [2], and attribution studies have linked anthropogenic warming to increased fuel aridity and forest-fire area in the western United States and California [3, 4]. These studies establish climatic influence at regional and seasonal scales. They do not, by themselves, tell us why one individual fire grows early and stops, another persists, and another waits before making its largest run. Addressing that question requires a response variable that retains the history.

Fire science has long recognized that much of an event's expansion can be concentrated in particular spread episodes, and weather-based definitions of spread-event days have been used to improve fire-growth modeling [5]. Comparisons of unusually large and smaller fires further show that growth responses can differ even when differences in daily weather are modest [6]. More recent analyses using

FIRED observations have related daily growth of individual western U.S. forest fires to aridity and prior fire size [7]. These studies motivate retaining the sequence of growth. Fire VASE contributes a complementary normalized geometry designed specifically for population-scale comparison of entire available observed histories.

Here we treat fire development as morphology. Fire VASE places relative developmental time on a common vertical axis and normalized cumulative burned area on a radial axis. Each observed developmental slice adds a ring; the mapped area increment determines how far the form widens. The result is both a visual object and a quantitative trajectory that preserves the timing, persistence, concentration, and recurrence of observed growth. The approach follows the general logic of morphometrics for quantifying form, functional data analysis for comparing trajectories, and principal component analysis for identifying dominant axes of variation [8, 9, 10]. The contribution is not the drawing alone. It is the creation of a common response geometry in which whole developmental histories can be compared.

Satellite-based approaches already reconstruct daily fire progression [11], track individual fires globally to derive size, duration, daily expansion, speed, and spread direction [12], and use subdaily observations to update perimeters and active fronts [13]. Such data reveal strong regional differences in fire duration and spread [14]. Fire VASE complements these spatial and event-metric approaches by turning the entire available observed sequence into a normalized geometry. The full archive contains 278,569 MODIS-derived FIRED events and 626,102 dated growth observations from 2000 to 2021 [15, 16, 17]. Because developmental morphology requires an internally ordered record, the primary coordinate system is fitted only to 10,246 events with at least three observations and no calendar-day gaps between the first and last retained observation. Daily gridMET weather is then projected onto that fire-derived representation for the 9,212 primary fires with complete weather [18].

We use the resulting space to ask a sequence of questions. Do observed histories vary along recurring developmental gradients? Does weather shift where fires occur in that space? Does current fire state alter the association between weather and subsequent observed growth? And what can be learned from cases in which weather and morphology disagree? Weather is the first demonstration of a broader explanatory architecture. Fire VASE makes developmental history the common response against which active-edge exposure, fuels, terrain, ignition context, suppression, and spatial progression can subsequently be tested.

2 Results

2.1 Dated growth observations become morphology

Figure 1 shows the transformation from dated burned-area increments to morphology. Relative developmental time runs upward, each ring marks an observed date, and normalized width tracks the square root of the fraction of reconstructed area accumulated. Abrupt widening records a large mapped addition of area; nearly parallel sides record successive dates with little additional mapped growth. Gray bands mark dates without an observation and are not treated as zero growth. Read from bottom to top, a VASE is a compressed history of how the fire accumulated area.

The VASE is not a geographic silhouette: width describes accumulated area, not perimeter shape or direction of travel. It is a standardized history. That distinction allows fires of different sizes and durations to be compared without erasing when their growth occurred.

Figure 1 illustrates the general encoding, including how gaps are displayed. Eligibility for the primary morphospace is stricter: only histories without intervening calendar-day gaps enter that coordinate system.

Validated v2 Figure 1

Add `Figure_1.png` from `figures/v2/`.

Open the validated repository figure.

Figure 1: **From dated growth observations to developmental morphology.** Four FIRED events illustrate the encoding. Upper panels show each observed daily share of reconstructed area as bars and cumulative share as lines. Gray bands denote unobserved dates, not zero-growth days. Lower panels encode the same records as VASEs: elapsed development runs upward, and ring radius is proportional to $\sqrt{A(t)/A(T)}$. The representation records normalized area allocation through time, not geographic perimeter shape or direction of spread.

2.2 Broad developmental gradients are robust to observation depth and geometry

The displayed VASE and the fitted morphospace are related but distinct. The VASE plots cumulative normalized area as radius; the PCA uses 20 normalized growth-allocation bins derived from that history. The primary fit includes 10,246 fires with at least three gap-free observations. Each bin records the share of observed growth assigned to one twentieth of relative developmental time, and the bins sum to one. Final area, duration, observation count, observed daily peak, mean growth, slenderness, and weather do not enter the fit. PC1 explains 34.1% of standardized variance, PC2 explains 28.2%, and the first five axes explain 89.4% (Fig. 2). Their leading loadings describe broad shifts in when growth occurs. Because the bins are proportions constrained to sum to one, these axes are a standardized Euclidean coordinate representation of compositional histories, not the only possible geometry for them.

A Hellinger sensitivity square-root transformed the allocation proportions and mean-centered them without variance scaling. It preserved the broad configuration: aligned PC1 and PC2 score correlations were 0.942 and 0.938, five-dimensional Procrustes similarity was 0.985, and pairwise-distance ranks correlated 0.976 with the primary space (fig. S4 and table S9). Local correspondence was weaker. Fifteen-neighbor overlap was 0.701 and overlap among extreme exemplars was 0.430.

The dominant early-versus-late allocation gradient remained recognizable, whereas the secondary concentrated-versus-distributed contrast rotated and weakened. The representation is therefore robust at the level of broad geometry, not pointwise invariant across analytical metrics.

These broad gradients persist as observation support becomes stricter. Refitting the representation to the 1,171 gap-free fires with at least seven observations produced correlations of 0.985 and 0.997 for the first two axes when the same reference fires were projected into both spaces. The rank correlation between five-dimensional pairwise distances was 0.969. Fine-scale structure was less stable: 15-neighbor overlap was 0.718, overlap among extreme exemplars was 0.606, and variance explained by the first five axes fell to 74.2%. Thus the major developmental directions are stable, whereas exact neighbors, extreme examples, and higher-dimensional detail depend more strongly on the observed record.

Temporal ordering showed systematic asymmetry relative to a within-fire permutation reference. In a 4,000-fire comparison, shuffling each fire's observed increments preserved its observation count, total area, and complete set of growth values while changing only their order. Observed histories allocated 0.541 of growth to the first half of development, compared with 0.500 after shuffling. They also had fewer detected pulses (1.249 versus 1.413) and reactivations (0.038 versus 0.120). These comparisons show that the specified order-sensitive summaries differ from this null; they do not establish a universal developmental program. Simulated positive growth allocations also show substantial low-dimensional compression.

Endpoint variables remain related to shape even though they do not define it. PC1 correlates with final area ($\rho = -0.308$), observed daily peak ($\rho = -0.332$), and duration or observation count ($\rho = -0.185$); PC2 correlates with final area ($\rho = 0.259$) and duration or count ($\rho = 0.256$). Developmental morphology is therefore constructed independently of endpoint attributes, not statistically independent of event scale or length. The coordinates are continuous, and interpretive landmarks aid description, but the existence or absence of latent natural fire classes was not tested.

Validated v2 Figure 2

Add `Figure_2.png` from `figures/v2/`.

Open the validated repository figure.

Figure 2: **The developmental morphospace is fitted from normalized growth allocation.** (A) Occupancy of the shape-only PC1–PC2 space for 10,246 fires with at least three observations and no intervening date gaps. (B) Loadings across the 20 relative-time allocation bins show how the first two axes redistribute normalized growth through development. (C) Observation support in the full archive distinguishes one-observation, two-observation, gappy, and primary gap-free histories. (D) Observed VASEs illustrate positions in the fitted space. PC1 explains 34.1%, PC2 28.2%, and the first five axes 89.4%; observation-threshold and temporal-order tests are reported in Fig. S2, and the Hellinger sensitivity in Fig. S4.

2.3 Weather maps weakly and heterogeneously onto morphology

The weather-complete primary population contains 9,212 fires. Event-mean VPD forms a diffuse gradient across the shape-only space rather than defining its axes (Fig. 3A). After adjustment for duration, observation count, area, region, month, and year, the highlighted associations remain small: mean VPD correlates negatively with late allocation ($\rho = -0.069$) and positively with front loading ($\rho = 0.050$), whereas precipitation correlates positively with late allocation ($\rho = 0.041$). The precipitation–pulse correlation falls from 0.070 to 0.025 after adjustment, and the interval obtained by resampling regions crosses zero.

Predictive performance likewise depends on the response (Fig. 3B). Among the tested ridge models, the displayed core weather means plus maximum VPD provide a compact common specification that can be compared across every response on the same cohort; they are not presented as a universally optimal predictor set. Region-held-out R^2 is 0.022 for front loading, 0.038 for late growth, 0.001 for peak timing, 0.019 for pulse count, 0.022 for reactivation, 0.074 for normalized entropy, and 0.080 for terminal taper. Normalized first and second differences are more predictable ($R^2 = 0.176$ and 0.161), but PCA coordinates recalculated within each validation fold are not. These weather summaries and linear models therefore weakly organize selected features of developmental morphology; they do not strongly recover whole histories in this representation.

This inference is geographically bounded. Although weather is complete for 237,235 archive events, the complete population is CONUS-only in these inputs; all Alaska and Hawaii events are excluded. Weather-complete histories are also longer and more frequently observed on average. These selection patterns, and the retrospective weather construction, limit generality beyond the 9,212-fire primary cohort.

Validated v2 Figure 3

Add `Figure_3.png` from `figures/v2/`.

Open the validated repository figure.

Figure 3: **External weather maps weakly and heterogeneously onto developmental morphology.** (A) Median event-mean VPD projected onto the shape-only morphospace for 9,212 weather-complete primary fires. Weather is external to the fitted axes. (B) Held-out R^2 for selected developmental responses using core event means plus maximum VPD under region, year, and strict crossed space–time holdouts. Predictive skill is small for timing and for shape coordinates recalculated within each validation fold. It is larger, though still limited, for entropy and normalized sequence differences. Adjusted associations and selection diagnostics are reported in the Supplementary Materials.

2.4 Recent fire state provides most subsequent-day predictive information

We next predicted growth on the exact subsequent calendar day using a common cohort of 87,944 transitions from 31,700 fires. As expected for a sequential growth process, the autoregressive state baseline—current growth, previous-day growth, cumulative observed area, and elapsed time—carried most transferable information. Its held-out R^2 was 0.448 under region and strict space–time holdouts and 0.458 under year holdout (Fig. 4A). The scientific question is therefore not whether recent growth predicts subsequent growth, but whether weather adds information after that baseline is known.

Adding day-specific weather increased R^2 by about 0.005. Adding all interactions between weather and fire state increased it by 0.015–0.018 above the autoregressive baseline (Fig. 4B). Interactions between VPD and current or previous growth contributed an additional 0.006–0.012 beyond the other interaction terms; the region-held-out increment was 0.0118. These average gains are reproducible, but their coefficients vary among regions, fire-size groups, and partial edge years. The result is a small amount of predictive information conditional on reconstructed fire state, not a universal VPD response curve or a causal weather effect.

Day-t weather was sampled at the centroid of that day's newly burned area, removing the use of final-event geometry in this comparison. The variables nevertheless remain retrospectively reconstructed end-of-day covariates. Satellite latency, date uncertainty, and time-zone alignment prevent interpreting the analysis as a prospective operational forecast.

Validated v2 Figure 4

Add `Figure_4.png` from `figures/v2/`.

Open the validated repository figure.

Figure 4: **Developmental state and exact subsequent-calendar-day growth.** (A) Total held-out R^2 for training-mean, current-growth persistence, autoregressive-state, day-t weather, additive state-plus-weather, and weather–state interaction models. (B) Paired held-out improvement above the shared autoregressive baseline; intervals resample fixed out-of-fold errors. (C) Day-t newly burned-area centroids and retrospective final-event centroids give similar small weather increments in the displayed sensitivity. (D) Seasonal estimates remain positive but heterogeneous. The shared cohort contains 87,944 transitions; these are retrospective associations, not a live forecast.

2.5 Matched contrasts provide diagnostic case studies

Caliper matching creates controlled candidates for follow-up by pairing only sufficiently similar fires. Within exact region, season, duration, and observation-count groups, and after applying limits on area ratio and standardized distance, 80.5% of eligible fires received a unique weather-space partner and 68.3% a unique morphology-space partner (Fig. 5A). Among each fire's ten nearest eligible neighbors, 91.2% had at least one acceptable weather match and 81.3% had at least one acceptable morphology match. These percentages varied with the chosen distance limit.

Mismatch was not greater than the conditional permutation reference. The fraction of weather-matched pairs with other-space distance greater than one was 49.7%, compared with a null mean of 50.4%; corresponding morphology-matched fractions were 39.5% and 40.2% (Fig. 5B). Thus the analysis does not establish unexpectedly frequent or ecologically prevalent mismatch.

The displayed pairs remain useful as controlled examples. One pair is close in normalized morphology but more separated in weather; the other is close in weather but follows contrasting cumulative growth paths (Fig. 5C,D). They are independent pairs selected near the median mismatch, not extreme anomalies or a two-by-two match. Fuel continuity, terrain, ignition context, suppression, active-edge exposure, and observation error remain hypotheses to test, not mechanisms inferred from these pairs [19, 20, 21, 22].

Validated v2 Figure 5

Add `Figure_5.png` from `figures/v2/`.

Open the validated repository figure.

Figure 5: **Representative matched-pair diagnostics.** (A) Fraction of 9,212 eligible fires assigned a unique acceptable partner within the same region, season, duration, and observation-count group, with an area ratio of at most 2 and a root-mean-square standardized distance of at most 0.5. (B) Observed fractions with distance greater than one in the other representation, compared with conditional within-group permutations; observed mismatch is compatible with the reference distribution. (C) A morphology-matched pair with differing weather. (D) An independently selected weather-matched pair with differing developmental histories. The pairs are diagnostic candidates, not extreme anomalies, a two-by-two match, or evidence identifying an omitted mechanism.

3 Discussion

Fire VASE changes the response being explained. The primary coordinates are fitted only from normalized allocation through relative developmental time, and the broad geometry survives removal of conventional endpoint variables and progressively stricter observation requirements. Within-fire permutation further shows systematic temporal asymmetry in front loading, pulsing, and reactivation under the specified null. Together, these results support developmental morphology as the paper's principal contribution.

The validation also sets its limits. Only 10,246 events—3.68% of the archive—have the gap-free observation support required for primary shape fitting, and no weighting extrapolates this selected cohort to the full archive. Retention varies strongly with event scale and record length: it is below 1% among fires no larger than 1 km^2, rises above 18% among fires of 10–100 km^2, and is highest for three-day histories before declining for longer records (Supplementary Table S8). The first two gradients and global distance structure remain recognizable across stricter observation thresholds and Hellinger geometry, but local neighborhoods, exemplar identity, secondary-axis interpretation, higher axes, and total five-axis coverage are less invariant. Low dimensionality is not uniquely diagnostic of biological restriction, and a continuous coordinate system neither proves nor disproves latent fire classes.

Excluding endpoints from coordinate construction clarifies rather than eliminates their relationship with development. Larger, longer, and faster-growing events remain moderately associated with shape coordinates. That association is plausible: event scale and developmental allocation arise from the same fire, even when they are represented separately. The analytical gain is that final area and duration can now be treated as external attributes to be explained alongside, rather than ingredients that define, developmental form.

Weather provides a first external test, not the main discovery. Its associations are small after coarse controls and depend strongly on the response being predicted. Weather recovers more of normalized roughness, taper, and entropy than precise peak timing or shape coordinates recalculated within each validation fold. The weakened precipitation–pulse result illustrates why adjustment and response-specific reporting matter. The data support modest external organization, not deterministic weather control of whole developmental histories.

The subsequent-day models reinforce this hierarchy. Recent fire state supplies the strong baseline; daily weather adds a small amount of transferable information, and interactions add somewhat more. Interactions involving VPD contribute reproducibly on average, but their coefficients vary too much to support a single response curve. These models show that environmental information depends on reconstructed fire state. They do not establish causation or operational availability.

Matched pairs make the framework useful for study design without implying that mismatch is more common than expected. The observed fractions agree with the conditional reference distribution, yet the balance requirements and distance limits still produce defensible candidate comparisons. The next step is to join developmental morphology with independently validated active-edge weather, fuel continuity, terrain, ignition context, suppression, and spatial progression. These layers can be tested against a response whose broad structure is already supported, rather than invoked after the fact to explain an unusual pair.

Important limitations remain. The strict cohort favors fires with complete consecutive observation histories and cannot be assumed to represent events excluded by observation depth or gaps. Satellite burn dates and gridMET days may not align perfectly, weather completeness excludes Alaska and Hawaii, and end-of-day fire geometry and exposure are reconstructed retrospectively. Uncertainty estimates treat the fitted predictions as fixed, regional resampling relies on only four broad regions, and predictive performance is weak for some small-fire groups. Pulse thresholds and matching distance limits are analytical choices. Within these boundaries, Fire VASE provides a robust normalized representation of how adequately observed fires developed, making it possible to ask not only why fires end differently, but why they take different paths to their endpoints.

4 Materials and Methods

The analysis proceeded from representation to explanation. We reconstructed dated FIRED growth histories, fitted a shape-only developmental coordinate system to adequately observed events, tested that representation against observation-depth and temporal-order nulls, projected weather as an external layer, and finally evaluated exact-day state models and matched diagnostic cases. Detailed definitions, sensitivities, complete validation grids, and provenance are reported in the Supplementary Materials.

4.1 Fire observations and primary population

Fire histories were derived from FIRED event and daily observation products based on MODIS burned area [15, 16, 17]. The analysis-ready dataset contains 278,569 events and 626,102 observations dated from 2 November 2000 through 1 May 2021. Every retained increment was joined by event identifier and date to the source daily geometry and checked against its mapped area. Histories with missing or nonnumeric values, negative increments, incomplete records, or zero reconstructed area were excluded rather than repaired or interpreted as low growth. The retained data contain no missing dates, duplicate dates, or zero-growth observations.

Of 347,533 adjacent observed pairs, 196,611 are exactly one calendar day apart, 81,940 are two days apart, and 68,982 span longer gaps. The full archive contains 161,073 one-observation events and 47,950 two-observation events. Among the 69,546 events with at least three observations, 59,300 contain at least one gap. The remaining 10,246 have no gaps between their first and last dated observation and define the primary developmental population. Gaps are unobserved dates, not imputed zero-growth days.

4.2 Normalized VASE geometry

For nonnegative daily increments g_j, reconstructed area was $S = \sum_j g_j$ and normalized growth allocation was $p_j = g_j/S$. For a positive-area fire, normalized VASE width at developmental time t was

$$w(t) = \sqrt{\frac{A(t)}{A(T)}}, \tag{1}$$

where $A(t)$ is cumulative reconstructed area and $A(T) = S$. Every VASE therefore terminates at width 1. Width is proportional, up to a constant, to the radius of an equal-area disk; it is not a geographic silhouette.

4.3 Shape-only developmental morphospace

In the primary consecutive cohort, each daily increment occupied one equal-width interval of relative developmental time. We interpolated cumulative allocation at 21 fixed boundaries and calculated the difference between adjacent boundaries, producing 20 nonnegative growth shares that sum to one. Interpolation standardizes the representation but does not create observations. The displayed VASE uses the cumulative allocation to determine radius, whereas these 20 incremental allocation bins were the only inputs to the primary PCA [10]. Final area, duration, observation count, true observed daily peak, mean growth, slenderness, scalar developmental traits, and weather were excluded. Each bin was centered and scaled using its training-set mean and standard deviation; near-constant columns received a scale factor of one. Because the shares form a composition, the resulting PCA is a standardized Euclidean coordinate representation rather than a unique geometry for the histories. Five axes were retained.

For the compositional sensitivity, allocations were transformed as $h_{ij} = \sqrt{p_{ij}}$. Each transformed row therefore had squared Euclidean norm one. Columns were mean centered but not variance scaled, preserving Hellinger distances up to a constant. PCA was then refitted to the same 10,246 fires. Components were matched by maximum absolute score correlation and oriented to positive aligned correlation. Global and local correspondence used the same Procrustes, score-subspace, pairwise-distance, nearest-neighbor, and exemplar-tail diagnostics as the observation-depth analysis.

Shape-oriented scalar traits were calculated after coordinate fitting. Front loading was allocation in the first half of relative development; late allocation was growth in the final quarter. Observed daily peak was the largest dated increment. Shannon entropy was $-\sum_j p_j \log p_j$, with zero-probability terms omitted, and normalized entropy divided by $\log n$ for $n > 1$ [23]. Normalized first and second differences described changes between successive allocation bins and were not interpreted as physical velocity or acceleration.

Interpretive landmarks were assigned by declared rules after excluding invalid, one-observation, two-observation, and gappy histories. Multiple detected pulses required at least two local maxima with prominence at least 20% of the observed peak. Reactivation required growth to return above 25% of peak after two consecutive subthreshold observed days. Late peak, front-loaded taper, and distributed growth used prespecified timing and allocation thresholds. These labels were not PCA inputs or statistically inferred natural classes.

4.4 Representation stability and temporal-order nulls

The unchanged 20-bin PCA was refitted for fires with at least 2, 3, 5, and 7 consecutive observations and for bootstrap, year, and region subsets. The same seeded set of 1,000 primary fires was projected into every fitted space, providing common reference points. After accounting for arbitrary component order and sign, we compared axis scores and loadings, five-dimensional subspaces, pairwise distances, nearest-neighbor sets, and the most extreme 2% of exemplars. These measures distinguish stability of the broad geometry from stability of local relationships.

Temporal-order comparisons used a recorded sample of 4,000 primary fires. Shuffling increments within each fire preserved observation count, total reconstructed area, and the complete set of observed increments. Symmetric Dirichlet allocations with concentration parameters of 1 and 10 provided alternative positive growth allocations with the same count and total. These simulations are reference conditions, not synthetic evidence. Reported tail probabilities are descriptive and were not adjusted for multiple comparisons.

4.5 Weather attribution and event-level models

After the fire-derived morphospace was constructed, daily gridMET weather was associated with each observed date [18]. Event-level morphology models summarized weather across the dated history. Exact-day state models instead projected that day's newly burned polygon, calculated its centroid, and sampled the nearest grid cell. These are retrospectively reconstructed end-of-day exposures rather than conditions proven available to a contemporaneous operational system.

The attributed variables were maximum and minimum temperature, VPD, wind speed, precipitation, relative and specific humidity, 100-hour and 1000-hour fuel moisture, energy release component, burning index, reference and potential evapotranspiration, and solar radiation. A history was complete only when every requested variable was available on every observed row. The archive contains 237,235 complete histories; the primary weather cohort contains 9,212.

Weather summaries included event means, quantiles, and fractions exceeding fixed thresholds. VPD exceedance was greater than 2 kPa, and a wet day had precipitation greater than 0.1 mm day^{-1}. Maxima, quantiles, and exceedance fractions were kept distinct because record length affects them differently. To calculate adjusted rank associations, ranked weather and response values were each regressed on duration, observation count, area, region, month, and year; the residuals were then correlated.

Event-level prediction used ridge regression with penalties 0.01, 1, and 100; penalty 1 was prespecified. Nested predictor sets separated core means, maximum VPD, comprehensive weather, and their combinations on the same 9,212 fires. The main figure displays core means plus maximum VPD as a compact common specification for comparing responses on an identical cohort, not as a universally optimal predictor set. Scaling used training data only. For shape PCs, PCA was refitted inside every training fold and held-out scores were centered consistently to prevent arbitrary fold offsets from inflating R^2.

Validation held out each coarse region, contiguous five-year blocks, or strict region-by-year-block intersections. In the crossed design, training excluded both the test region and test year block. Random-fire folds were retained only as an interpolation diagnostic. Held-out R^2 used a training-derived reference mean and retained negative values. Fire-, year-, and region-cluster intervals used 200 resamples of fixed out-of-fold errors and therefore do not include model-refitting or remote-sensing uncertainty.

4.6 Exact subsequent-day growth models

State models used only transitions separated by exactly one calendar day and required an observation on the previous day for every comparison. The common cohort contained 87,944 transitions from 31,700 fires. The response was $\log[1 + g_{t+1}(\mathrm{km}^2)]$. The autoregressive state baseline contained current growth, previous-day growth, cumulative reconstructed area, and elapsed time. Models then added current-day weather and interactions between weather and state. All models used the same observations, training-only scaling, the prespecified ridge penalty, and the same validation blocks as the event-level models.

VPD-specific validation removed or added the interactions of VPD with current and previous growth while retaining the other interaction terms. Paired bootstrap intervals tested changes in held-out performance using fixed predictions. Coefficients expressed in the original units used standard errors clustered by fire and remained conditional on the fitted design and penalty. Refits by region, year, season, fire size, and training fold assessed heterogeneity. Tables of VPD by fire state reported both transitions and distinct fires without extrapolating a response surface beyond the observed data.

4.7 Matched diagnostic pairs

Matching used either core event-mean weather variables or the 20 morphology bins, standardized within the 9,212-fire cohort. Potential partners were the 10 nearest neighbors within the same region, season, duration, and observation-count group. Pairs with an area ratio greater than 2 or a root-mean-square standardized distance greater than 0.5 were rejected. The remaining pairs were sorted by distance and identifiers and selected in order without reusing a fire; differences in the other representation did not influence selection. Sensitivity analyses varied the number of candidate neighbors, distance measure, and maximum acceptable distance. Conditional reference distributions shuffled the other representation within the same adjustment groups and an additional area class, $\lfloor \log_2(\text{area}) \rfloor$.

4.8 Reproducibility and scope

All analyses use the analysis-ready v0.1 FIRED/gridMET data lake, with no synthetic fallback. The validated baseline SHA is `18c923cb0c82bf9f66567b62a3491ac30a28c369`; the freeze manifest records subsequent code and output hashes. Configuration `config/analysis_v2.json` uses seed 20260828. Checks of scientific invariants, reference-data conservation, manuscript claims, publication hashes, deterministic regeneration, and figure copies passed across 66 outputs. Dedicated validation tests pass. The broader test collection retains a disclosed pre-existing error in two modules that import a missing helper.

Daily gridMET and satellite burn dates have timing and spatial limitations. Nearest-cell exposures omit within-fire heterogeneity, directional wind, fuel continuity, terrain, suppression, and active-edge conditions. The primary and weather populations are selected by observation support and completeness. All model and matching results are observational and retrospective.

References and Notes

References

[1] J. K. Balch *et al.*, The fastest-growing and most destructive fires in the US (2001 to 2020). *Science* **386**, 425–431 (2024). doi:10.1126/science.adk5737.

[2] A. L. Westerling, H. G. Hidalgo, D. R. Cayan, T. W. Swetnam, Warming and earlier spring increase western U.S. forest wildfire activity. *Science* **313**, 940–943 (2006). doi:10.1126/science.1128834.

[3] J. T. Abatzoglou, A. P. Williams, Impact of anthropogenic climate change on wildfire across western US forests. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 11770–11775 (2016). doi:10.1073/pnas.1607171113.

[4] A. P. Williams, J. T. Abatzoglou, A. Gershunov, J. Guzman-Morales, D. A. Bishop, J. K. Balch, D. P. Lettenmaier, Observed impacts of anthropogenic climate change on wildfire in California. *Earth's Future* **7**, 892–910 (2019). doi:10.1029/2019EF001210.

[5] J. Podur, B. M. Wotton, Defining fire spread event days for fire-growth modelling. *Int. J. Wildland Fire* **20**, 497–507 (2011). doi:10.1071/WF09001.

[6] B. E. Potter, D. McEvoy, Weather factors associated with extremely large fires and fire growth days. *Earth Interact.* **25**, 160–176 (2021). doi:10.1175/EI-D-21-0008.1.

[7] C. S. Juang, A. P. Williams, J. T. Abatzoglou, J. K. Balch, M. D. Hurteau, M. A. Moritz, Rapid growth of large forest fires drives the exponential response of annual forest-fire area

to aridity in the western United States. *Geophys. Res. Lett.* **49**, e2021GL097131 (2022). doi:10.1029/2021GL097131.

[8] F. L. Bookstein, *Morphometric Tools for Landmark Data: Geometry and Biology* (Cambridge Univ. Press, 1991).

[9] J. O. Ramsay, B. W. Silverman, *Functional Data Analysis*, ed. 2 (Springer, 2005). doi:10.1007/b98888.

[10] I. T. Jolliffe, *Principal Component Analysis*, ed. 2 (Springer, 2002). doi:10.1007/b98835.

[11] S. Veraverbeke, F. Sedano, S. J. Hook, J. T. Randerson, Y. Jin, B. M. Rogers, Mapping the daily progression of large wildland fires using MODIS active fire data. *Int. J. Wildland Fire* **23**, 655–667 (2014). doi:10.1071/WF13015.

[12] N. Andela, D. C. Morton, L. Giglio, R. Paugam, Y. Chen, S. Hantson, G. R. van der Werf, J. T. Randerson, The Global Fire Atlas of individual fire size, duration, speed and direction. *Earth Syst. Sci. Data* **11**, 529–552 (2019). doi:10.5194/essd-11-529-2019.

[13] Y. Chen *et al.*, California wildfire spread derived using VIIRS satellite observations and an object-based tracking system. *Sci. Data* **9**, 249 (2022). doi:10.1038/s41597-022-01343-0.

[14] J. A. Balik, J. D. Coop, M. A. Krawchuk, C. E. Naficy, M.-A. Parisien, S. A. Parks, C. S. Stevens-Rumann, E. Whitman, Biogeographic patterns of daily wildfire spread and extremes across North America. *Front. For. Glob. Change* **7**, 1355361 (2024). doi:10.3389/ffgc.2024.1355361.

[15] L. Giglio, L. Boschetti, D. P. Roy, M. L. Humber, C. O. Justice, The Collection 6 MODIS burned area mapping algorithm and product. *Remote Sens. Environ.* **217**, 72–85 (2018). doi:10.1016/j.rse.2018.08.005.

[16] J. K. Balch, L. A. St. Denis, A. L. Mahood, N. P. Mietkiewicz, T. M. Williams, J. McGlinchy, M. C. Cook, FIRED (Fire Events Delineation): An open, flexible algorithm and database of US fire events derived from the MODIS burned area product (2001–2019). *Remote Sens.* **12**, 3498 (2020). doi:10.3390/rs12213498.

[17] A. L. Mahood, E. J. Lindrooth, M. C. Cook, J. K. Balch, Country-level fire perimeter datasets (2001–2021). *Sci. Data* **9**, 458 (2022). doi:10.1038/s41597-022-01572-3.

[18] J. T. Abatzoglou, Development of gridded surface meteorological data for ecological applications and modelling. *Int. J. Climatol.* **33**, 121–131 (2013). doi:10.1002/joc.3413.

[19] J. K. Balch *et al.*, Human-started wildfires expand the fire niche across the United States. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 2946–2951 (2017). doi:10.1073/pnas.1617394114.

[20] M.-A. Parisien, M. A. Moritz, Environmental controls on the distribution of wildfire at multiple spatial scales. *Ecol. Monogr.* **79**, 127–154 (2009). doi:10.1890/07-1289.1.

[21] L. M. Holsinger, S. A. Parks, C. Miller, Weather, fuels, and topography impede wildland fire spread in western US landscapes. *For. Ecol. Manage.* **380**, 59–69 (2016). doi:10.1016/j.foreco.2016.08.035.

[22] N. A. Povak, P. F. Hessburg, R. B. Salter, Evidence for scale-dependent topographic controls on wildfire spread. *Ecosphere* **9**, e02443 (2018). doi:10.1002/ecs2.2443.

[23] C. E. Shannon, A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423 (1948). doi:10.1002/j.1538-7305.1948.tb01338.x.

Acknowledgments

Funding: Funding information to be added before submission.

Author contributions: Author contributions to be completed before submission.

Competing interests: The authors declare no competing interests.

Data and materials availability: The external FIRED, MODIS burned-area, and gridMET inputs are publicly available from the cited sources. Derived analysis tables, figure-generation scripts, manuscript-generation code, and the shareable Fire VASE data-lake package are organized in the Fire VASE repository. Repository DOI, CyVerse accession, or archival accession to be added before submission.

AI transparency: OpenAI Codex/ChatGPT was used as an AI-assisted coding, analysis, visualization, and editorial tool during development of this project. AI assistance included drafting and revising Python scripts for Fire VASE data ingestion, climate attribution, morphospace analysis, statistical summaries, figure generation, PDF/report production, and render-based quality checks; drafting and revising manuscript text, figure legends, response-to-review material, and simulated reviewer critiques; searching for and organizing candidate citations and author-guideline requirements; and helping maintain logs, manifests, tests, schemas, and documentation.

The AI system did not originate the underlying FIRED, MODIS burned-area, gridMET, PRISM, or other observational data, did not make final scientific judgments independently, and is not listed as an author. The author directed the analyses, selected the scientific claims, reviewed the code and reported outputs, and remains responsible for the integrity, interpretation, citations, and final content of the manuscript. Independent domain review remains necessary before submission.

Synthetic or illustrative demonstrations created during repository development are documented separately and were not used as evidentiary data for the manuscript analyses.

Supplementary Materials

Supplementary Methods S1 to S9, Supplementary Results S10 to S14, Figs. S1 to S4, and Tables S1 to S9 are provided in a separate Supplementary Materials document.