

Seasonal Landsat summaries improve cross-fire prediction of post-fire conifer establishment

Joshua Puliafico

JoshuaPulifico63296@outlook.com

Independent Researcher <https://orcid.org/0009-0001-2602-6466>

Research Article

Keywords: post-fire regeneration, Landsat, grouped validation, average precision, remote sensing, forest monitoring

Posted Date: September 22nd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11100941/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Seasonal Landsat summaries improve cross-fire prediction of post-fire conifer establishment

Joshua Puliafico^{a,*}

^a*Independent Researcher, Cuba*

Abstract

Satellite information is useful for post-fire monitoring only if it remains informative when predictions are transferred to new wildfire events. We compared an ecological and survey-context baseline (M0), a seasonally summarized Landsat representation (M1), and a larger representation containing additional 32-day within-season descriptors (M2). Models were trained on 25 fires and evaluated on 13 prespecified held-out fires from the same western U.S. field compilation. In the primary Year 1–3, 90 m analysis, M1 increased fire-weighted average precision (AP) relative to M0 by 0.1180 (95% fire-bootstrap interval 0.0738–0.1632) for Elastic Net and by 0.1031 (0.0350–0.1701) for histogram gradient boosting. The M1 improvement was positive in 12 of 13 held-out fires for both model families. M2 did not produce a further held-out improvement. A post hoc prefire-only control also exceeded M0, so the full Landsat increment cannot be assigned specifically to post-fire recovery. Across all 38 fires, a development-informed reduced representation (Mc; 214 predictors) performed similarly to M1 (350 predictors), with Mc–M1 differences of –0.0009 and +0.0046. Gains varied among contributor and geographic domains. At a 20% within-fire assessment capacity, M1 increased the mean proportion of selected plots without observed conifers from 0.622 to 0.729 for Elastic Net and from 0.549 to 0.630 for boosting. Under the tested early observation window, seasonal Landsat summaries improved cross-fire ranking and field-sample prioritization, whereas denser within-season descriptors did not add a corresponding gain.

Keywords: post-fire regeneration, Landsat, grouped validation, average precision, remote sensing, forest monitoring

1. Study rationale and objectives

Establishment of new conifers is a central step in post-fire forest recovery. Seed availability, burn severity and water stress can limit establishment even where mature conifers were present before the fire, thereby influencing future forest composition and reforestation needs (Stevens-Rumann et al., 2018; Davis et al., 2019; Coop et al., 2020). Regional field compilations have clarified several of these constraints, including the short-term buffering associated with lower fire severity (Davis et al., 2023). For management, however, the practical challenge is to obtain comparable information across many fire events and use it to focus limited field-survey effort (Fargione et al., 2021).

Landsat offers a long record of vegetation conditions before and after disturbance. Greenness, near-infrared response and moisture-sensitive indices capture related but distinct properties of the land surface (Tucker, 1979; Gao, 1996; Jiang et al., 2008; Badgley et al., 2017; Bright et al., 2019). These signals are not specific to conifer recruitment: shrubs, herbs, hardwood resprouting and surviving canopy can all contribute to apparent spectral recovery (Vanderhoof et al., 2021; Celebrezze et al., 2024). Recent synthesis work also shows that recovery metrics and field measurements can describe different trajectories and associations (Smith-Tripp et al., 2026). The central question is therefore predictive rather than purely descriptive: does a satellite representation improve prediction of a defined field outcome beyond a meaningful ecological baseline?

The value of temporal detail is less clear. Seasonal aggregation may preserve persistent differences in vegetation condition while reducing sensitivity to cloud gaps, short-lived signals and local timing variation. More detailed descriptors of amplitude, green-up, decline, peaks and harmonics could add ecological information, but they also increase predictor correlation and dependence on local phenology. A representation that performs well on development fires may therefore transfer poorly to unseen events. Evaluation must consequently hold out complete fires; random plot-level splits would allow shared fire context to leak across the training and test sets (Roberts et al., 2017).

We tested a prespecified sequence of predictor representations: ecological and survey context, seasonal Landsat summaries, and additional within-season descriptors. The primary experiment trained models on 25 fires and evaluated them on 13 held-out fires from the same field compilation. Elastic Net and histogram gradient boosting provided sparse linear and nonlinear model families. We first asked whether seasonal summaries improved held-out discrimination and then whether greater temporal detail added further improvement. Later analyses examined a reduced three-index candidate, domain heterogeneity, prefire information and the use of predictions for field-sample prioritization. This sequence distinguishes the original held-out comparison from development-informed exploratory analyses. Figure 1 presents the study workflow and separates the primary Pilot25-to-Locked13 experiment from the later All38 benchmark.

2. Study design and analytical workflow

2.1. Field response definition and cohort construction

The analysis used the field compilation released by Davis et al. (2023), which records conifer regeneration density together

*Corresponding author.

Email address: JoshuaPuliafico63296@outlook.com
(JoshuaPuliafico)




Figure 1. Study overview and transfer logic. The illustration summarizes the sequence from post-fire landscapes and repeated Landsat observations to seasonal feature construction, event-level validation and field-screening outputs. M0 denotes contextual information, M1 seasonal summaries, M2 additional within-season descriptors and Mc the compact exploratory set. This is a conceptual diagram rather than an observed landscape image or causal model. The artwork was generated with GPT Image 2 and reviewed by the author; all quantitative results come from the archived analyses.

with site conditions and survey metadata. The original table contained 10,230 plots. After checks of coordinates, timing and fire identity, 8,364 plots associated with 291 source fires remained. We defined the response as field-detected conifer establishment: plots with total conifer density greater than zero were coded as present, and plots with density equal to zero as absent. Missing density values were retained as missing rather than recoded as absence. This endpoint represents establishment observed during the field survey; it is not a measure of adequate stocking, species-specific suitability or eventual stand recovery.

We reconciled field fire names and years with Monitoring Trends in Burn Severity boundaries using fire year, spatial containment, distance and name agreement (Eidenshink et al., 2007). The crosswalk preserved original event keys, aliases and decisions about multipart fires. The primary spatial assignment required a stable wildfire identity and a plot located within the matched boundary; records near the boundary edge were reserved for sensitivity analyses.

We first assembled a 55-fire candidate pool using sample support, representation of both response classes and diversity criteria. Screening for Landsat support retained 38 fires with the required Year 1–3 observations at both spatial supports. Pilot25 contained 25 development fires and 2,114 plots, whereas Locked13 contained the other 13 eligible fires and 1,241 plots. Together, All38 contained 3,355 plots (Table 1; Fig. 2). Because response support contributed to the earlier candidate-pool construction, these fires are not a probability sample of all western U.S. fires. Locked13 is an event-level held-out cohort from the same compilation, not an independently collected external field program.

Table 1. Cohort composition and analytical role of each dataset.

	Development	Held-out	All fires
Fires	25	13	38
Sampled plots	2,114	1,241	3,355
Observed presence	1,002	620	1,622
Observed absence	1,112	621	1,733
Median survey age	13 (5–21)	5 (5–15)	Union of cohorts
Analytical role	Model development	Event-level test	Exploratory benchmark

Values are reported at plot level unless otherwise noted. Survey age is the plot-level median with range. The combined set is the union of the two cohorts, not a third external sample.

2.2. Satellite observation windows and preprocessing

We processed Landsat Collection 2 Level 2 surface-reflectance observations in Google Earth Engine (Gorelick et al., 2017). The extraction workflow allowed Landsat 5, 7, 8 and 9, although the contributing sensors differed among fire-specific observation periods. Optical digital numbers were converted as $\text{reflectance} = \text{DN} \times 0.0000275 - 0.2$. We masked fill values, clouds, cloud shadow, snow and radiometric saturation, while retaining sensor-composition and Landsat 7 gap diagnostics. No empirical cross-sensor radiometric harmonization was applied.

Observations were assigned to nine consecutive 32-day bins from 1 March through 30 November; the final bin was truncated at the seasonal boundary, so the bins are not calendar months. The primary window included three complete prefire years (–3 to –1) and post-fire years +1 to +3, excluding the fire year. Where survey timing allowed, years +1 to +5 were used in secondary analyses. Images had to precede 1 January of the survey year. Empty bins were retained as explicit missing observations rather than being interpolated into phenological dates.

At 30 m, extraction used the native centre pixel. The pri-

mary 90 m support was obtained by taking the median across a 3×3 native-pixel neighbourhood for each image before temporal summarization. The retained indices were NDVI, EVI2, NBR, NDMI and NIRv. The frozen exports contained 293,472 point-scale-bin rows for Pilot25 and 138,636 for Locked13; 6.55% and 10.63% of the respective complete exports were missing rows. Figure S4 compares missingness over the shared six-year primary window.

2.3. Predictor blocks and temporal availability

M0 comprised ten groups of contextual predictors: relativized burn ratio; annual and growing-season climatic water deficit; heat-insolation load; distance to the nearest live seed source; surrounding post-fire live-tree cover within 300 m; years from fire to survey; plot area; prefire disturbance category; and contributor. The climatic-deficit variables are 1981–2010 normals rather than weather measurements collected progressively after each fire. Fire identifiers, coordinates, batch identifiers and any outcome-derived fields were excluded.

M1 added 340 seasonally aggregated predictors to M0. For each index and year, quality-controlled bins were summarized by location, spread, extrema, an observed seasonal integral and measures of observation support. Prefire baselines and dispersion were used to form absolute or standardized anomalies, recovery fractions and slopes. M1 therefore combines antecedent spectral context, post-fire observations and information about data availability; it should not be interpreted as a pure intervention on post-fire recovery. In the archive, ANNUAL denotes this March–November seasonal aggregate, not a full calendar year.

M2 added 600 within-season predictors to M1, including seasonal means, green-up and decline summaries, persistence, amplitude, peak and trough timing, and first- and second-order harmonic terms. Harmonic fits required at least six valid bins spanning spring, summer and autumn. The raw inventories contained 10 columns for M0, 350 for M1 and 950 for M2, before categorical encoding; these counts are not estimates of effective degrees of freedom or computation time. A prefire-only control added seasonal descriptors from years -3 to -1 to M0.

This is a retrospective prediction study based on field outcomes and early Landsat windows. Several baseline variables, including survey age and field-derived seed-source information, do not by themselves define an operational forecast available at post-fire year 3. We examined reduced-context refits, but did not establish a common year-3 availability timestamp for every baseline variable. A deployable early-warning system would require that additional availability audit and a fixed forecast horizon.

2.4. Event-level validation and model comparison

We fitted Elastic Net logistic regression and histogram gradient boosting (Zou and Hastie, 2005; Friedman, 2001; Pedregosa et al., 2011). The primary specification used the Year 1–3 window at 90 m. Within Pilot25, five repeats of five-fold stratified group cross-validation held out complete fires. Four-fold grouped validation within each outer training set selected hyperparameters. Imputation, scaling, categorical encoding and other transformations were fitted only on training data. Candidate models used the same outer partitions and sample mask.

After development, final pipelines were fitted using all Pilot25 fires and applied to Locked13. The original held-out

comparison was kept separate from feature definition, preprocessing, fitting, tuning, calibration and model-family selection. Subsequent exploratory use of Locked13 neither changed the archived predictions nor created a second untouched test set.

2.5. Performance estimands and uncertainty assessment

The primary discrimination measure was fire-weighted average precision (AP), called AUCPR in the archive. We used the non-interpolated average-precision definition rather than trapezoidal integration of the precision–recall curve. For plot i in fire f , with n_f evaluated plots, the weight was $w_i = 1/n_f$. Precision and recall were calculated from the pooled weighted predictions, so each fire contributed the same total weight. This estimand is not the arithmetic average of separate fire-level AP values. The weighted outcome prevalence provides the reference precision shown in the precision–recall plots.

For repeated grouped validation, the five out-of-fold probabilities were first averaged within plot and AP was then calculated. Each plot therefore contributes once to the reported point estimate. Pilot25 summaries based on concatenated repeat-level predictions are retained only as an explicitly labelled archival comparison in Supplementary Table S1. Locked13 uses one set of frozen final-pipeline predictions. Paired differences are calculated as $AP(A) - AP(B)$ on matching plots, not as differences between bootstrap means.

The original 95% fire-bootstrap intervals are based on 2,000 paired resamples of complete fires. We retained the archived endpoints and predictions, while revising the tables to distinguish observed point differences from bootstrap averages. These intervals describe variation across the sampled fires conditional on the fitted pipelines. They do not include all uncertainty from feature-family selection or repetition of the development workflow. Later diagnostics use their own resampling procedures and are identified as exploratory.


2.6. Reduced representations and transfer diagnostics

A true-refit feature-family analysis in Pilot25 identified aggregated NDVI, EVI2 and NIRv predictors as a candidate reduced family. The resulting Mc specification combined M0 with 204 predictors from those indices, for 214 columns in total. Mc was fixed before evaluation on All38, using modal hyperparameter choices from earlier Pilot25 refits and the existing grouped partitions. No additional reduced subsets were searched at that stage.

The Mc choice was informed by earlier Pilot25 outcomes, and Pilot25 is included in All38. The selected feature family and settings were therefore not independent of every All38 fold. We interpret this comparison as an exploratory reduced-representation benchmark rather than a fully nested or independently confirmed test. Its intervals do not account for the upstream selection, and similar performance with intervals spanning zero is not formal equivalence because no equivalence margin was specified.

For contributor-transfer refits, we used Pilot25 and removed every fire containing the held-out contributor. Contributor-specific test sets can overlap because a fire may contain plots from several contributors. Figure 5a–b reports all ten comparisons with their test-fire and test-plot denominators; six contributors have at least three held-out fires. The complete archived contrast matrix appears in Fig. S5. We also evaluated five fixed geographic fire blocks over All38. These post

Pilot25: 25 fires / 2,114 plots | Locked13: 13 fires / 1,241 plots | Observed sample support



EN, Elastic Net; HGB, histogram gradient boosting.

Figure 2. Cohort coverage and observation support. Panels a–b locate the fires and summarize climatic and burn-severity context. Panels c–e show survey timing, outcome prevalence and sampled-plot counts. Panel f displays missing observation fractions across the shared relative-year window, and panel g standardizes fire-level covariates using the development cohort. The figure describes the analytical sample and measurement support; it is not a map of regional coverage or evidence of independent plot-level replication.

hoc analyses retained earlier-selected hyperparameters and describe domain sensitivity rather than independent, newly tuned external validation.

2.7. Robustness checks and information-source controls

Archived refits successively removed contributor, contributor plus plot area, and then seed-source distance and surrounding tree cover. We also examined common survey-age support and eligibility exclusions for the source *Pinus contorta* model. Deleting one Locked13 fire at a time from the frozen predictions measured test-set influence without refitting. Supplementary Methods documents the warning-free All38 checkpoint repair and computational provenance.

We examined the complexity result with true-refit feature-family contrasts, lower-dimensional seasonal additions and fixed-partition alternative representations. The phase-invariant sensitivity removed explicit timing and phase descriptors; it did not align observations to local snowmelt or phenological time. The sensor-control sensitivity added identical sensor-composition and missingness summaries to M1 and M2, but was not a radiometric harmonization.

We also made a post hoc paired comparison between M1 and the frozen prefire-only control on Locked13. Complete fires were resampled 2,000 times with seed 20260906, using the same resampled fires for both models. Parameters and predictions were not changed. Because the two representations differ in construction and dimensionality, this comparison diagnoses the specificity of the Landsat increment; it does not identify a causal post-fire effect.

2.8. Field-screening utility and recovery summaries

To connect prediction with field prioritization, we added an absence-directed screening diagnostic. Within each Locked13 fire, plots were ranked from lowest to highest predicted establishment probability. We evaluated fixed assessment capacities of 10%, 20%, 30% and 50%, selecting the ceiling of each fraction of the fire's plot count and using plot identifiers to break ties. Precision was the proportion of selected plots without observed conifers, averaged equally across fires. This quantity differs from presence-oriented AP and from a universal probability threshold. Marginal and paired 95% intervals used 2,000 complete-fire bootstrap resamples with seed 20260905. The diagnostic measures ranking enrichment among archived sample plots, not intervention efficacy or landscape-scale treatment benefit.

Reliability plots use ten fixed probability bins and the same fire-equal weights. We summarized the existing Pilot25-trained Platt-calibrated probabilities without recalibration. For descriptive spectral context, we calculated outcome-specific median index values within each fire and then summarized those values across fires by their median and interquartile range. The prefire point represents years -3 to -1 . These trajectories are descriptive, not model effect estimates, and their interquartile ranges are not confidence intervals.

2.9. Heterogeneity, model agreement and visualization protocol

For each evaluated fire, we calculated AP, area under the receiver-operating-characteristic curve (AUROC) and Brier score in addition to the pooled fire-weighted metrics. Within-fire paired differences indicate where gains occurred, but their sign counts are descriptive and do not replace the pooled AP

estimand or whole-fire bootstrap. Brier score was the weighted mean squared error between predicted probability and observed presence; a positive reduction therefore indicates lower probability error. All comparisons used matching plots and frozen probabilities (Fig. 3c–h).

For Mc and M1 in All38, we calculated the Spearman rank correlation and mean absolute probability difference across all 3,355 plots, and then recalculated rank correlation within each fire. We also scored predictions from each of the five archived outer cross-validation repetitions to describe sensitivity to partitioning. Every repetition contained the same 38 fires and 3,355 plots, so they are not independent datasets. Averaging their probabilities within plot reproduced the stored mean-out-of-fold predictions to numerical precision (Fig. 4f–i).

For the spectral diagnostics, we first calculated outcome-group medians within each fire. Across fires, we then summarized presence-minus-absence differences and the number of fires with each direction for all five indices. Cross-index correlations used fire-level, outcome-group recovery changes between years $+1$ and $+3$. This paired construction separates within-fire contrasts from changes caused by the composition of fires (Fig. 9). Observation-completeness matrices used a common six-year window, while descriptive standardized covariates used Pilot25 fire-level means and standard deviations (Fig. 2f–g).

Versioned Python scripts generated the numerical figures from archived predictions, source observations and explicitly derived tables. ChatGPT (OpenAI; GPT-6 Astra Pro for the present assembly) assisted with manuscript organization, code preparation and consistency checks; the reproducibility package contains the executed computations, source tables and figure-level provenance. No model was refitted and no new bootstrap sample was drawn during this figure-integration stage. Figure 1 is now a conceptual illustration generated with GPT Image 2; its landscapes, satellite and map elements are explanatory artwork rather than observed research images.

3. Findings

3.1. The evaluation fires covered distinct temporal and geographic settings

Pilot25 included 1,002 presence and 1,112 absence plots, whereas Locked13 included 620 and 621, respectively (Table 1). Median survey age was 13 years in Pilot25 and 5 years in Locked13. The centroid distributions and ecological covariate support also differed (Fig. 2), confirming that the held-out comparison involved transfer between fire events rather than interchangeable plots. All 13 held-out fires contributed to the primary analysis, but only three supported the Year 1–5 sensitivity. Those secondary-window results therefore do not describe the same fire population (Fig. S3).

3.2. Seasonal summaries improved ranking on unseen fires

Using one mean out-of-fold probability per plot, Pilot25 AP rose from 0.6803 to 0.7027 for Elastic Net and from 0.6343 to 0.6814 for boosting when M1 replaced M0. The held-out comparison was more informative. In Locked13, AP increased from 0.5855 to 0.7035 for Elastic Net and from 0.5743 to 0.6774 for boosting (Fig. 3; Table 2). The observed M1–M0 differences were $+0.1180$ (95% fire-bootstrap interval 0.0738–0.1632) and $+0.1031$ (0.0350–0.1701), against a weighted held-out prevalence of 0.5287. Both model families therefore

gained discrimination from the aggregated Landsat representation. Within-fire AP differences were positive in 12 of 13 held-out fires for both families (Fig. 3c). Fire-weighted Brier scores also improved, falling from 0.2660 to 0.2342 for Elastic Net and from 0.2882 to 0.2658 for boosting (Fig. 3h). M2–M1 differences were positive in only six and seven fires, respectively (Fig. 3d).

Held-out AP for M2 was 0.6923 for Elastic Net and 0.6554 for boosting. Relative to M1, the observed differences were -0.0112 (-0.0409 to 0.0123) and -0.0219 (-0.0414 to 0.0069). Both intervals included zero. Thus, the added within-season descriptors did not demonstrate an improvement in this held-out experiment, although the result does not show that every possible temporal-detail representation would have exactly zero benefit. The corresponding development mean-OOF differences were -0.0475 and -0.0164 .

Table 2. Held-out performance for the primary Year 1–3, 90 m comparison.

Set / learner	M0 context	M1 seasonal	M2 detail
Pilot25			
Elastic Net	0.6803	0.7027	0.6552
Pilot25			
HGB	0.6343	0.6814	0.6650
Locked13			
Elastic Net	0.5855	0.7035	0.6923
Locked13			
HGB	0.5743	0.6774	0.6554

Set / learner	M1 minus M0	M2 minus M1
Pilot25		
Elastic Net	+0.0224 [-0.0178 , $+0.0601$]	-0.0475 [-0.0748 , -0.0172]
Pilot25		
HGB	+0.0470 [-0.0124 , $+0.0942$]	-0.0164 [-0.0592 , $+0.0301$]
Locked13		
Elastic Net	+0.1180 [$+0.0738$, $+0.1632$]	-0.0112 [-0.0409 , $+0.0123$]
Locked13		
HGB	+0.1031 [$+0.0350$, $+0.1701$]	-0.0219 [-0.0414 , $+0.0069$]

Paired differences; 95% fire-bootstrap intervals in brackets.

Pilot25 AP is computed from one mean out-of-fold probability per plot. Locked13 uses frozen final-pipeline predictions. Intervals are the archived 95% paired fire-bootstrap intervals; points are observed differences, not bootstrap means.

3.3. The compact predictor set approached the exploratory performance plateau

Across All38, Elastic Net AP was 0.6504 for M0, 0.7081 for Mc, 0.7090 for M1 and 0.6875 for M2. The corresponding boosting values were 0.6397, 0.7065, 0.7019 and 0.6797 (Fig. 4; Table 3). Mc exceeded M0 by $+0.0577$ (0.0224 – 0.0938) for Elastic Net and $+0.0668$ (0.0261 – 0.1023) for boosting. Compared with M1, Mc differed by -0.0009 (-0.0080 to 0.0076) and $+0.0046$ (-0.0090 to 0.0173), respectively.

Mc used 136 fewer raw columns than M1 and showed similar average discrimination in this exploratory comparison. It did not establish superiority or formal equivalence, and its selection history prevents All38 from being treated as an independent confirmation. M2 remained below M1 in All38, with observed differences of -0.0216 and -0.0222 and intervals below zero. The results are consistent with a plateau among the tested representations, while the original Locked13 experiment remains the strongest held-out evidence. Mc and M1

probabilities were highly rank-correlated over All38 (Spearman coefficients 0.9673 and 0.9461), but their mean absolute differences were 0.0407 and 0.0555 (Fig. 4f–g). Similar AP therefore did not mean identical plot-level probabilities. Within-fire agreement and the five archived repetition-specific contrasts provide complementary views of that variation (Fig. 4h–i).

Table 3. Exploratory comparison of the compact representation across all eligible fires.

Learner	Context	Compact set	Seasonal set	Detailed set
Elastic Net	0.6504	0.7081	0.7090	0.6875
HGB	0.6397	0.7065	0.7019	0.6797

Learner	Mc minus M0	Mc minus M1
Elastic Net	+0.0577 [$+0.0224$, $+0.0938$]	-0.0009 [-0.0080 , $+0.0076$]
HGB	+0.0668 [$+0.0261$, $+0.1023$]	+0.0046 [-0.0090 , $+0.0173$]

Paired differences; 95% fire-bootstrap intervals in brackets.

Mc selection used earlier Pilot25 outcomes. The All38 benchmark is not fully selection-nested, and its intervals do not account for that upstream choice.

3.4. Transfer performance varied across source and geographic domains

Transfer varied across the ten contributor-held-out comparisons, with the corresponding fire and plot denominators retained for every contrast (Fig. 5a–b). Among the six contributors represented by at least three test fires, Elastic Net Mc–M0 was positive for two and negative for four, while M1–M0 was positive for one and negative for five. Boosting Mc–M0 was positive for four of six. Because contributor test sets can overlap, these contrasts describe source-domain sensitivity rather than six independent replications.


Geographic blocks showed the same lack of uniformity. In block 4, Mc reduced but did not remove the negative increment: Elastic Net differences were -0.0177 for Mc–M0 and -0.0340 for M1–M0, while boosting differences were -0.0235 and -0.0513 . Thus, an average cross-fire gain coexisted with domains in which the ecological baseline performed better. The reduced representation did not by itself resolve geographic or survey-program transfer.

3.5. Context restrictions and leave-one-fire checks preserved the main pattern


The M1 advantage remained after removing contributor, contributor plus plot area, and the broader restricted-context set (Fig. 6). Across these three archived refits, M1–M0 differences ranged from $+0.1058$ to $+0.1202$ for Elastic Net and from $+0.0822$ to $+0.1422$ for boosting; the corresponding fire-bootstrap intervals remained above zero. M2–M1 differences were small or negative with intervals that included zero. The no-survey-design boosting estimate was slightly positive ($+0.0022$), so the added-detail result was not uniformly negative.

Deleting each of the 13 held-out fires in turn from the frozen predictions retained M1 above M0. The smallest increments were $+0.1018$ for Elastic Net and $+0.0721$ for boosting. The largest M2–M1 increments remained negative (-0.0007 and -0.0137). These deletions quantify the influence of individual evaluated fires. Analyses restricted by survey-age common


a Elastic Net · PR




b Gradient boosting · PR




c Aggregate gain by fire




d Added-detail effect by fire




e Pooled paired effects



f EN · error trade-off



g HGB · error trade-off




h Absolute held-out performance

Family / metric	M_0	M_1	M_2	Family / metric	M_0	M_1	M_2
EN · AP	0.5855	0.7035	0.6923	HGB · AP	0.5743	0.6774	0.6554
EN · AUROC	0.5739	0.6762	0.6638	HGB · AUROC	0.5507	0.6228	0.6121
EN · Brier	0.2660	0.2342	0.2468	HGB · Brier	0.2882	0.2658	0.3098


EN, Elastic Net; HGB, histogram gradient boosting.

Figure 3. Performance transfer to held-out fires. Panels a–b show fire-weighted precision–recall curves. Panels c–d give the paired AP change for each test fire, while panel e summarizes the pooled contrasts with archived whole-fire bootstrap intervals. Panels f–g relate AP improvement to Brier-score change, and panel h compares AP, AUROC and Brier score on the same held-out predictions. Repeated development predictions are averaged within plot before scoring; fire-specific points are descriptive rather than separate confidence intervals.

a Discrimination vs complexity



b Probability error vs complexity




c Input modules


	Ctx	Veg	Wet	Detail
M_0	10	0	0	0
M_c	10	204	0	0
M_1	10	204	136	0
M_2	10	204	136	600

Counts, not importance


d Pooled AP contrasts




e Heterogeneity across 38 fires




f EN · M_c versus M_1




g HGB · M_c versus M_1



h Within-fire rank agreement




i Five previously fitted outer-CV repetitions




EN, Elastic Net; HGB, histogram gradient boosting.

Figure 4. Feature-size trade-offs and prediction stability. The panels compare raw predictor counts, feature-family composition, paired AP changes, plot-level probability agreement and repetition-specific results for M_0 , M_c , M_1 and M_2 . Hexagonal density plots summarize the M_c – M_1 probability relationship across All38, and the lower panels show within-fire agreement and the five archived outer-validation repetitions. The compact set was selected from earlier development analyses and is therefore treated as exploratory.


a EN · source-held-out refits




b HGB · source-held-out refits




c Fixed geographic blocks



d EN · block holdouts



e HGB · block holdouts




f Block outcome support


Block	Fires	Plots	Prev.	Age
B1	15	1701	0.52	5.0
B2	3	174	0.66	17.0
B3	9	472	0.52	14.0
B4	7	700	0.44	12.0
B5	4	308	0.32	13.5

Prev.: equal-fire prevalence
Age: median fire-median survey age

g Source · learner agreement



h Geography · learner agreement



EN, Elastic Net; HGB, histogram gradient boosting.

Figure 5. Variation across source and geographic domains. Panels a–b show contributor-held-out contrasts with the corresponding test-fire and test-plot denominators. Panel c maps the five fixed geographic blocks; panels d–e report block-level paired differences, and panel f summarizes block composition. Panels g–h compare model-family agreement across domains. These are archived transfer diagnostics with different training and testing populations, not independent replication studies.

support or eligibility also retained $M1 > M0$ and $M2 < M1$ (Fig. S3).

3.6. Ranking models improved prioritization under limited field capacity

At a 20% assessment capacity within each held-out fire, each model selected 255 plots. The mean absence proportion rose from 0.6225 with M0 to 0.7294 with M1 for Elastic Net, and from 0.5487 to 0.6301 for boosting (Fig. 7). Paired improvements were +0.1069 (95% fire-bootstrap interval 0.0528–0.1639) and +0.0814 (0.0217–0.1477). M2 did not provide a clear additional gain: M2–M1 differences were -0.0062 (-0.0493 to 0.0323) and $+0.0083$ (-0.0596 to 0.0769). The fire-mean absence prevalence across all test plots was 0.4713. Capacity-specific recall and the complete 13-fire precision-difference matrix show the associated coverage–enrichment trade-off and local heterogeneity (Fig. 7c–f).

These are post hoc enrichment results among observed plots, not estimates of restoration effects. The raw reliability curves departed from the identity line, particularly at extreme predictions (Fig. 7). Existing Pilot25-trained calibration reduced Brier scores in the archived sensitivity (Fig. S2), but neither calibration result nor ranking analysis establishes a universal threshold that can be transported to new domains.


3.7. Prefire information qualifies a recovery-specific interpretation

True-refit feature-family comparisons produced a larger development gain for the aggregated vegetation-index family than for the moisture/burn-index family. Additional morphology or harmonic descriptors did not establish a stable improvement beyond M1 (Fig. 8). Lower-dimensional seasonal controls, phase-invariant comparisons and sensor-composition checks did not overturn the early-window complexity pattern (Figs. S4 and S5).


The frozen prefire-only control reached Locked13 AP values of 0.6578 for Elastic Net and 0.6566 for boosting. Its post hoc differences from M0 were +0.0724 (0.0132–0.1381) and +0.0822 (0.0024–0.1547). Differences between M1 and the prefire-only control were +0.0456 (-0.0049 to 0.1104) and +0.0208 (-0.0273 to 0.0636). These positive point estimates were not separated from zero by the fire-bootstrap intervals. The complete M1 gain over M0 therefore cannot be attributed uniquely to newly observed post-fire recovery.

Both outcome groups showed post-fire spectral change across all five indices, with interquartile ranges that overlapped among fires (Fig. 9a–e). At year +3, the within-fire NDVI median was higher for presence than absence plots in 30 of 38 fires, and the corresponding NBR count was 28 of 38. Similar contrasts were already common before fire: 29 of 38 fires for NDVI and 31 of 38 for NBR (Fig. 9f–h). These patterns are compatible with a contribution from antecedent spectral context. They provide context for the prediction results but do not identify conifer-specific reflectance, seedling-emergence timing or a causal recovery pathway.


a EN · context refits




b HGB · context refits




c Delete one test fire: aggregate gain




d Delete one test fire: added detail



e Temporal and spatial support




f Survey-age and eligibility restrictions




EN, Elastic Net; HGB, histogram gradient boosting.

Figure 6. Sensitivity to context, fire influence and observation support. Panels a–b summarize context-removal refits, panels c–d show leave-one-fire changes from frozen predictions, and panel e compares temporal windows and spatial supports. Panel f contrasts the primary, common-age and restricted-eligibility test sets. The panels use the supplied uncertainty information; no interval is added where the archive did not provide one.


a EN · precision of absence screening




b HGB · precision of absence screening




c EN · absence recall




d HGB · absence recall




e Fire-specific screening gain at 20%




f Mean screening gain at 20%



g EN · calibration



h HGB · calibration




EN, Elastic Net; HGB, histogram gradient boosting.

Figure 7. Utility for prioritizing field assessment. Panels a–d quantify precision and absence recall when only a fixed fraction of plots is assessed. Panel e gives fire-specific changes at 20% capacity, panel f summarizes the paired contrasts, and panels g–h show reliability by fixed probability bins. These results describe enrichment among sampled plots and should not be interpreted as a landscape-wide treatment effect.


a True-refit feature inventory

ID	Ctx	Agg.	Detail	Total
A0	10	0	0	10
A1	10	204	0	214
A2	10	136	0	146
A3	10	340	0	350
A4	10	340	300	650
A5	10	340	250	600
A6	10	340	600	950


b Paired refit effects on Pilot25




c Compressed seasonal additions




d Archived prefire-control bootstrap



e Prefire contrasts across test fires








f Bounded alternative explanations





EN, Elastic Net; HGB, histogram gradient boosting.

Figure 8. Information-source and dimensionality checks. The figure contrasts the archived predictor inventories, refit effects, low-dimensional seasonal additions, prefire-control bootstrap distributions, observed within-fire contrasts and fixed-partition sensitivities. The prefire and M1 representations differ in construction, so the comparison is diagnostic rather than a causal decomposition of Landsat information. No new model fitting or resampling was performed for this figure.

a NDVI · observed trajectories

b EVI2 · observed trajectories

c NIRv · observed trajectories

d NBR · observed trajectories

e NDMI · observed trajectories

f Direction consistency across fires

NDVI	29/38	33/38	30/38	30/38
EVI2	19/38	22/38	21/38	23/38
NIRv	20/38	21/38	20/38	21/38
NBR	31/38	34/38	31/38	28/38
NDMI	29/38	33/38	28/38	26/38
	Pre	+1	+2	+3

Presence median > absence median

g NDVI · outcome contrast

h NBR · outcome contrast

i Cross-index recovery concordance

NDVI	1.00	0.90	0.87	0.75	0.78
EVI2	0.90	1.00	0.98	0.84	0.85
NIRv	0.87	0.98	1.00	0.81	0.83
NBR	0.75	0.84	0.81	1.00	0.97
NDMI	0.78	0.85	0.83	0.97	1.00
	NDVI	EVI2	NIRv	NBR	NDMI

Spearman ρ across 38 fires
Recovery: year +3 minus year +1
Outcome groups weighted equally

EN, Elastic Net; HGB, histogram gradient boosting.

Figure 9. Spectral trajectories and cross-index agreement. Panels a–e show fire-level and across-fire summaries for five indices. Panel f counts the direction of within-fire presence–absence contrasts, panels g–h show selected index differences in a stable fire order, and panel i reports correlations among recovery contrasts. The trajectories summarize observed groups and do not represent species-specific growth, learned feature importance or a causal pathway.

4. Interpretation and limitations

4.1. Why seasonal Landsat summaries transferred across fires

The primary held-out experiment indicates that the tested seasonal Landsat representation improved discrimination of later field-observed conifer establishment beyond ecological and survey context. The pattern was consistent across two model families and remained after several context-variable removals. The main contribution is therefore a controlled comparison of information representations on complete unseen fires, not a claim that spectral greening directly measures conifer recruitment.

This interpretation is consistent with field-linked remote-sensing studies showing that non-conifer vegetation can dominate growing-season recovery signals (Vanderhoof et al., 2021; Celebrezze et al., 2024), and that spectral and field recovery depend on what is measured and when (Smith-Tripp et al., 2026). In the present analysis, prefire spectral information already improved prediction, while the additional M1-versus-prefire-control contrast remained uncertain. M1 may combine antecedent vegetation condition, persistent site properties, post-fire trajectories and observation support. Separating those components would require a matched prefire-versus-postfire feature design.

4.2. Additional temporal detail did not yield a consistent benefit

Within the primary Year 1–3 window, 600 additional seasonal descriptors did not demonstrate a held-out gain over the aggregated summaries. The comparison included both sparse linear and nonlinear models, and smaller seasonal additions were also examined. The result favours preserving information that transfers across fires rather than increasing the number of transformations applied to an already informative time series.

Redundancy, non-conifer vegetation and context-dependent timing are plausible explanations, but these experiments do not establish a mechanism. Removing phase-related descriptors reduced part of the performance deficit without aligning observations to ecological time. Sensor-composition controls likewise did not remove sensor-specific radiometric differences. Snow-covered observations, structural measurements or locally aligned phenological time may contain different information (Vanderhoof and Hawbaker, 2018; Menick et al., 2024). The current conclusion is limited to the tested calendar-based early-window representations; the three-fire Year 1–5 sensitivity cannot support an equally broad claim about longer windows.

4.3. A smaller predictor set is a practical follow-up benchmark

The exploratory complexity comparison places the 214-column Mc candidate close to the 350-column seasonal model. Mc is therefore a practical candidate for a future locked comparison. The reduction concerns the raw predictor inventory only; computational savings, missing-data robustness and equivalent predictive utility were not independently benchmarked here.

The selection chronology remains important. Mc was fixed before the final All38 benchmark, but earlier outcomes from fires retained in that benchmark motivated its construction. Fixing a development-informed subset does not make its evaluation independent of development. A rigorous follow-up

should nest index-family selection inside every training partition or apply the fixed candidate to a genuinely new field program. That design would show how much of the apparent plateau survives independent feature-family selection.

4.4. The clearest application is targeted field screening

The absence-directed diagnostic translates ranking performance into a field-assessment question: among a fixed fraction of sampled plots in each fire, how often did the lowest-probability group lack observed conifers? At the illustrated 20% capacity, M1 improved this enrichment over M0. That result connects the prediction task more directly to reconnaissance than a presence-oriented metric alone, while leaving restoration decisions to local ecological and silvicultural assessment.

Several operational boundaries remain. The response records any observed establishment, not adequate density or desired species composition. Plot areas and survey ages vary, and some baseline predictors depend on field or survey information. Calibration is not sufficiently stable to justify a universal probability cutoff. A deployable monitoring product would need a defined information-availability date, common forecast horizon, locally relevant endpoint and prospective testing. The present results support research on field prioritization, not a validated prescription to plant or withhold treatment.

4.5. Inference remains bounded by the sampled fire domains

Locked13 supplies genuine held-out fire events, but development and evaluation still come from the same source compilation. Earlier support screening selected fires with sufficient sample size, outcome variation and imagery. Contributor and regional analyses show that the mean gain is not uniform across domains. The 13 held-out fires also provide a modest event denominator despite the larger number of plots. The archived fire-bootstrap intervals condition on the fitted pipelines and do not capture every component of model-selection uncertainty.

The most informative next evidence would combine independent regional or contributor-level replication, a matched information-source comparison and a fixed-horizon availability audit. These designs address different questions: transportability, the incremental value of post-fire observations beyond prefire context, and whether the model is usable at the claimed decision time. Keeping these questions separate preserves the value of the current cross-fire result without extending it beyond what the data support.

5. Concluding remarks

Seasonal Landsat summaries improved held-out prediction of field-observed conifer establishment beyond the ecological and survey-context baseline in the tested Year 1–3, 90 m analysis. Adding within-season descriptors did not demonstrate a further held-out gain. Because a prefire-only control also improved on the baseline, the complete Landsat increment cannot be interpreted as uniquely representing post-fire recovery. The reduced NDVI/EVI2/NIRv candidate performed similarly to the full seasonal representation in the exploratory All38 analysis, but its selection was development-informed and requires an independently selected comparison. Overall, the results

support seasonal Landsat information as a component of cross-fire field-prioritization models while emphasizing domain heterogeneity, predictor timing and calibration as requirements for the next validation stage.

Data and analytical resources

Field observations originate from the public archive associated with Davis et al. (2023); Landsat and MTBS products are public. The frozen reproducibility package accompanying this submission includes the source items used in the analysis, feature matrices, grouped partitions, archived predictions, inherited fitting scripts, diagnostic source tables, plotting code and integrity manifests. The current manuscript is linked to those sources through a figure-number crosswalk and an evidence manifest. The package distinguishes rapid numerical verification and figure regeneration from the optional full model-refitting workflow. The original providers' terms continue to apply to their materials; a permanent public accession and the licences for original code and derived tables will be supplied with the publication archive.

Author contribution statement

JoshuaPuliafico: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Resources, Visualization, Project administration, Writing – original draft, Writing – review and editing.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of competing interest

The author declares no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We acknowledge the investigators and field crews whose observations were compiled by Davis et al. (2023), and the agencies maintaining the Landsat and MTBS archives.

Use of generative AI and AI-assisted tools

ChatGPT (OpenAI; GPT-6 Astra Pro for the present assembly) assisted language revision, manuscript organization, programming and numerical-figure integration. GPT Image 2 (OpenAI; model gpt-image-2) was used to create the conceptual artwork for Figure 1. That artwork is an explanatory illustration, not an observed field image or a spatial prediction product. The numerical results and statistical figures were derived from the archived analysis outputs and documented computational workflow; they were not generated by an image model. The author is responsible for checking the AI-assisted material, confirming its originality and accuracy, and approving the final scientific interpretation.

References

- Badgley, G., Field, C.B., Berry, J.A., 2017. Canopy near-infrared reflectance and terrestrial photosynthesis. *Science Advances* 3, e1602244. <https://doi.org/10.1126/sciadv.1602244>.
- Bright, B.C., Hudak, A.T., Kennedy, R.E., Braaten, J.D., Henareh Khalyani, A., 2019. Examining post-fire vegetation recovery with Landsat time series analysis in three western North American forest types. *Fire Ecology* 15, 8. <https://doi.org/10.1186/s42408-018-0021-9>.
- Celebrezze, J.V., et al., 2024. A fast spectral recovery does not necessarily indicate post-fire forest recovery. *Fire Ecology* 20, 54. <https://doi.org/10.1186/s42408-024-00288-6>.
- Coop, J.D., Parks, S.A., Stevens-Rumann, C.S., Crausbay, S.D., Higuera, P.E., Hurteau, M.D., Tepley, A., Whitman, E., Assal, T., Collins, B.M., Davis, K.T., Dobrowski, S., Falk, D.A., Fornwalt, P.J., Fulé, P.Z., Harvey, B.J., Kane, V.R., Littlefield, C.E., Margolis, E.Q., North, M., Parisien, M.-A., Prichard, S., Rodman, K.C., 2020. Wildfire-driven forest conversion in western North American landscapes. *BioScience* 70, 659–673. <https://doi.org/10.1093/biosci/biaa061>.
- Davis, K.T., Dobrowski, S.Z., Higuera, P.E., Holden, Z.A., Veblen, T.T., Rother, M.T., Parks, S.A., Sala, A., Maneta, M.P., 2019. Wildfires and climate change push low-elevation forests across a critical climate threshold for tree regeneration. *Proceedings of the National Academy of Sciences* 116, 6193–6198. <https://doi.org/10.1073/pnas.1815107116>.
- Davis, K.T., et al., 2023. Reduced fire severity offers near-term buffer to climate-driven declines in conifer resilience across the western United States. *Proceedings of the National Academy of Sciences* 120, e2208120120. <https://doi.org/10.1073/pnas.2208120120>.
- Eidenshink, J., Schwind, B., Brewer, K., Zhu, Z.-L., Quayle, B., Howard, S., 2007. A project for monitoring trends in burn severity. *Fire Ecology* 3, 3–21. <https://doi.org/10.4996/fireecology.0301003>.
- Fargione, J., Haase, D.L., Burney, O.T., Kildisheva, O.A., Edge, G., Cook-Patton, S.C., Chapman, T., Rempel, A., Hurteau, M.D., Davis, K.T., Dobrowski, S., Enebak, S., De La Torre, R., Bhuta, A.A.R., Cubbage, F., Kittler, B., Zhang, D., Guldin, R.W., 2021. Challenges to the reforestation pipeline in the United States. *Frontiers in Forests and Global Change* 4, 629198. <https://doi.org/10.3389/ffgc.2021.629198>.
- Friedman, J.H., 2001. Greedy function approximation: A gradient boosting machine. *Annals of Statistics* 29, 1189–1232.
- Gao, B.-C., 1996. NDWI - A normalized difference water index for remote sensing of vegetation liquid water from space. *Remote Sensing of Environment* 58, 257–266. [https://doi.org/10.1016/S0034-4257\(96\)00067-3](https://doi.org/10.1016/S0034-4257(96)00067-3).
- Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., Moore, R., 2017. Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment* 202, 18–27. <https://doi.org/10.1016/j.rse.2017.06.031>.
- Jiang, Z., Huete, A.R., Didan, K., Miura, T., 2008. Development of a two-band enhanced vegetation index without a blue band. *Remote Sensing of Environment* 112, 3833–3845. <https://doi.org/10.1016/j.rse.2008.06.006>.
- Menick, C., Tinkham, W., Hoffman, C., Vanderhoof, M., Vogeler, J., 2024. Snow-cover remote sensing of conifer tree recovery in high-severity burn patches. *Remote Sensing of Environment* 305, 114114. <https://doi.org/10.1016/j.rse.2024.114114>.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830.
- Roberts, D.R., Bahn, V., Ciuti, S., Boyce, M.S., Elith, J., Guillerma-Arroita, G., Hauenstein, S., Lahoz-Monfort, J.J., Schröder, B., Thuiller, W., Warton, D.I., Wintle, B.A., Hartig, F., Dormann, C.F., 2017. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography* 40, 913–929. <https://doi.org/10.1111/ecog.02881>.
- Smith-Tripp, S.M., Coops, N.C., Meyer, J., White, J.C., Gergel, S., Axelsson, J., Roeser, D., 2026. Synthesizing Spectral and Field Observations of Post-fire Conifer Recovery in Dry Conifer Forests. *Ecosystems* 29, 11. <https://doi.org/10.1007/s10021-025-01029-9>.
- Stevens-Rumann, C.S., Kemp, K.B., Higuera, P.E., Harvey, B.J., Rother, M.T., Donato, D.C., Morgan, P., Veblen, T.T., 2018. Evidence for declining forest resilience to wildfires under climate change. *Ecology Letters* 21, 243–252. <https://doi.org/10.1111/ele.12889>.
- Tucker, C.J., 1979. Red and photographic infrared linear combinations

- for monitoring vegetation. *Remote Sensing of Environment* 8, 127-150. [https://doi.org/10.1016/0034-4257\(79\)90013-0](https://doi.org/10.1016/0034-4257(79)90013-0).
- Vanderhoof, M.K., Hawbaker, T.J., 2018. It matters when you measure it: Using snow-cover Normalised Difference Vegetation Index (NDVI) to isolate post-fire conifer regeneration. *International Journal of Wildland Fire* 27, 815-830. <https://doi.org/10.1071/WF18075>.
- Vanderhoof, M.K., Hawbaker, T.J., Ku, A., Merriam, K., Berryman, E., Cattau, M., 2021. Tracking rates of postfire conifer regeneration vs. deciduous vegetation recovery across the western United States. *Ecological Applications* 31, e02237. <https://doi.org/10.1002/eap.2237>.
- Zou, H., Hastie, T., 2005. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B* 67, 301-320. <https://doi.org/10.1111/j.1467-9868.2005.00503.x>.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SUPPLEMENTARYMATERIAL.pdf](#)