Leakage-audited machine learning versus ETAS for earthquake forecasting in the Sea of Marmara

Authors

Basri Kerem Alhan
Independent researcher
Email: basrikerem@gmail.com

Kenessary Khabat
Independent researcher
Email: khabatkenessary@gmail.com

Preprint status

This manuscript is a non-peer-reviewed preprint submitted to EarthArXiv. It has not yet undergone formal peer review and has not yet been accepted for publication in a journal.

Corresponding author

Basri Kerem Alhan
basrikerem@gmail.com

Leakage-audited machine learning versus ETAS for earthquake forecasting in the Sea of Marmara

Basri Kerem Alhan and **Kenessary Khabat**

Corresponding author: Basri Kerem Alhan (basrikerem@gmail.com)

Code & data: https://github.com/keremalhan/marmara-forecast

Abstract

The Sea of Marmara hosts a seismic gap directly beneath a metropolitan region of some 18 million people, and the 23 April 2025 Mw 6.2 Kumburgaz earthquake renewed attention on whether short-term forecasting can add value there. We report a leakage-audited forecasting experiment built on a strictly causal, homogenized KOERI catalogue (31,329 model-box events, 2003–2026). Motivated by the finding that machine-learning models rarely beat the Epidemic-Type Aftershock Sequence (ETAS) model, we designed the benchmark around a machine-checkable causality gate — a truncated-catalogue self-test that recomputes every feature from a catalogue truncated to the forecast time and requires exact reproduction. Within this design we (i) replaced first-generation ETAS with a conditional cascade Monte-Carlo forecaster, which ranks best on M≥3.5/30-day targets (precision–recall AUC 0.130 vs 0.126 for first-generation ETAS) and produces roughly 1.8× more offspring inside active sequences; (ii) built an ETAS×ML hybrid that by construction cannot lose to the cascade but does not beat a properly-fit ETAS in information gain (−0.38 nats per event); (iii) trained a conditional "bigger-event-ahead" discriminator on renewal-timed synthetic sequences, which ranks the escalating 2025 Sındırgı sequence far above the decaying Kumburgaz sequence; and (iv) introduced an information-arrival analysis — a general procedure that freezes the forecast at successive lead times and measures when, and through which observable, a target event first becomes visible above the background hazard. Applied to the 2025 Marmara events, it shows that spatial localization is available far in advance (the 23 April Mw 6.2 fault cell was the top ~1% seismicity cell throughout the preceding year) while the specific timing is bounded by foreshock physics: that event's cell was quiescent, and only the lone M4.5 foreshock lifted the 30-day M≥6 probability gain to 42×, ten minutes before the mainshock. We report all negative results (machine learning ≤ ETAS on rare targets, a null short-term aftershock-incompleteness correction, no measurable gain from static GNSS strain, and a spatial-coverage miss for the out-of-box Sındırgı sequence). The system issues probability *gain*, not alarms; the live 30-day regional M≥6 probability at 2026-07-05 is a fraction of a percent, the same order as the Poisson base rate.

Keywords: earthquake forecasting; ETAS; machine learning; data leakage; Sea of Marmara; operational earthquake forecasting; foreshocks; information gain

1. Introduction

The Main Marmara Fault, the submarine continuation of the North Anatolian Fault beneath the Sea of Marmara, is the most prominent unbroken segment of a fault system that ruptured progressively westward across the twentieth century in a sequence of large earthquakes that culminated in the 1999 Mw 7.4 İzmit and Mw 7.2 Düzce events to its east. The central portion of the Marmara fault has not produced a major earthquake since 1766, and it lies directly offshore of Istanbul, a metropolitan region of some 18 million people. Long-term renewal calculations place the decade-scale probability of an M ≥ 7 beneath the sea in the tens of percent (Parsons, 2004), and the fault's persistent locking leaves it a mature seismic gap. Recent work has sharpened this picture: an analysis of nearly two decades of seismicity documents an eastward progression of Mw > 5 events along the fault over the past ~15 years, breaching formerly quiet segments — affecting both creeping and locked sections — in what the authors describe as a steady march toward the city (Martínez-Garzón et al., 2025). The 23 April 2025 Mw 6.2 (Mw 6.3 in some solutions) Kumburgaz earthquake, a strike-slip event at ~6 km depth on the central Marmara segment that triggered several hundred aftershocks within two weeks, is the most recent step in that sequence (Ertuncay et al., 2025). The combination of a mature seismic gap, a very large exposed population, and an active, migrating sequence makes short-term forecasting for this region both consequential and scientifically demanding — and makes it especially important that any reported "skill" be robust to the evaluation artifacts described below rather than an artifact of them.

The operational standard for short-term forecasting is ETAS (Ogata, 1988), which models seismicity as a background rate plus self-exciting Omori–Utsu aftershock cascades. ETAS underpins the operational aftershock forecasting now run by national agencies (Hardebeck et al., 2024; Mizrahi et al., 2024). A wave of neural point-process and deep-learning models has been proposed as a more flexible alternative (Dascher-Cousineau et al., 2023), but the most careful head-to-head benchmark to date, EarthquakeNPP, found that neural models do **not** systematically beat a well-implemented ETAS when the two are compared under consistent, reproducible conditions (Stockman et al., 2025). This is the empirical backdrop for the present study: a machine-learning forecasting experiment for Marmara must account for the ETAS baseline and guard against the evaluation artifacts that inflate apparent machine-learning skill.

Chief among those artifacts is **data leakage** — the accidental flow of information from the evaluation target back into the features or the model-selection procedure. Leakage is endemic in machine-learning studies of seismicity. Common forms include oversampling positives by row duplication *before* a train/test split (so copies of test positives sit in training), tuning a decision threshold or a kernel half-life on the test set, ablating features with future information visible rather than forward-chained, and a "leakage test" that never actually permutes the target and so cannot detect leakage at all. Because rare large events are the quantity of interest, even a small leak of target information is enough to manufacture apparent M6 skill that is really base-rate rescaling, and such skill routinely evaporates to precision = recall = 0 under forward validation. We therefore adopt an **audit-first** design and, critically, make the absence of look-ahead *machine-checkable* rather than merely asserted: a truncated-catalogue self-test recomputes every feature from a catalogue truncated to the forecast time and requires bit-for-bit reproduction, so any reader can verify causality directly rather than trusting an assertion of it.

This paper makes six contributions. First, a **leakage-audited benchmark** for Marmara, whose causality is enforced by a truncated-catalogue self-test that any reproducer can run. Second, a **conditional cascade Monte-Carlo forecaster** that replaces first-generation ETAS `expected_counts` with full recursive, spatially anisotropic simulation. Third, a **structurally non-inferior ETAS×ML hybrid**, which is competitive with but not superior to ETAS. Fourth, a **renewal-informed synthetic training scheme** for a conditional large-event discriminator, because the real catalogue contains too few large events to train on directly. Fifth, an **information-arrival analysis** — a general method that freezes the forecast at successive lead times to quantify when, and through which observable, a target event becomes visible above background — which we apply to the 2025 Marmara events and which decomposes the forecasting problem into a solved spatial part and a foreshock-bounded temporal part. Sixth, a running **hashed, append-only prospective protocol**, so that future forecasts accrue genuine out-of-sample credibility that cannot be gamed by hindsight. Throughout, we report probability *gain* relative to base rates and retain the negative results.

2. Data and study region

Region and grid. The model box spans 25.6–30.9° E and 39.6–41.9° N, chosen to cover the Main Marmara Fault system and the northern strands of the North Anatolian Fault where the KOERI network is densest, so that completeness is as low and as uniform as the data allow (Martínez-Garzón et al., 2025). Forecasts are made on a 0.1° grid (1,219 model-box cells) in 30-day windows; the 0.1° spacing (~8–11 km at these latitudes) is a compromise between fault-scale spatial resolution and having enough events per cell for stable feature estimation, and the 30-day window matches the horizon at which first-generation aftershock forecasting is reliable. For rare-event training only, we additionally use a wide box (25.0–31.5° E, 39.0–42.5° N, 2,275 cells) that captures more M≥4.5 events — including the 2011 Simav, 2013 Aegean, 2014 Gökçeada, 2022 Düzce, and 2025 Sındırgı events — without changing the evaluation region; models trained on the wide box are always evaluated on the model box alone. The Sındırgı sequence lies south of the model box and is therefore used only as an out-of-box coverage check, a limitation we return to in the results.

Catalogue. We derive a homogenized, deduplicated catalogue from the KOERI regional bulletin, 2003-01-04 to 2026-07-06, comprising 31,329 events in the model box . Three data-level faults in the source were corrected before any modelling. (1) Timestamps in the bulletin are in Türkiye local time but had been treated as UTC; they were converted to true UTC and verified empirically by matching a well-recorded anchor event's origin time to its independently reported UTC time, fixing the offset for the whole catalogue (the proof is recorded in `data/fetch_manifest.json`). An uncorrected three-hour timestamp error would silently corrupt every time-since-last-event and short-window count feature. (2) A single, untyped magnitude column pooled ML, Md, Mw, and Ms without conversion; magnitudes were typed and homogenized to a proxy-moment magnitude (`mag_w`) using the empirical relations of Kadirioğlu and Kartal (2016), applied — and flagged — slightly below their stated validity ranges for a small tail. Because the ML→Mw and Md→Mw offsets differ, this typing is what makes the completeness mixture (below) tractable and keeps the M≥3.5 and M≥4.5 targets on a single, consistent scale. (3) 11,754 rows identified as quarry blasts were screened out on the basis of time-of-day,

location, and discrimination criteria; counting anthropogenic blasts as tectonic seismicity would bias both the background rate and the feature counts. The audit that motivated these corrections is documented in the repository (`docs/AUDIT.md`).

Completeness. The max-curvature magnitude of completeness is Mc = 3.65 (mode 3.45 + 0.2; Wiemer & Wyss, 2000) . This value is *inflated* by ~15,000 Md events piling near `mag_w` 3.45, because the Md→Mw offset is steeper than ML→Mw; the dominant modern ML population itself completes near `mag_w` 2.72. Because our targets are defined on `mag_w`, the M≥3.5 target (y35) remains above ML-population completeness and the M≥4.5 target (y45) is far above Mc. Mc is applied identically to feature counting and to every baseline (fair). The mixed-catalogue character makes the Gutenberg–Richter b-value the dominant source of M6 extrapolation uncertainty — and the estimated b is itself sensitive to catalogue-processing choices such as declustering, which can shift the mainshock size distribution materially (Mizrahi et al., 2021) — so we carry a b-value ensemble: Aki (1965) maximum-likelihood b ≈ 1.02, a b-positive estimate ≈ 1.54 (van der Elst, 2021), and a calibration-consistent operational b_op = 1.2 (real-window regression slope ≈ 1.04). Model-box counts are 5,021 events ≥ 3.5, 192 ≥ 4.5, 12 ≥ 5.5, and a single M ≥ 6 (the 2025 Mw 6.2) . The single real M6 is why every M6 number in this study comes from Gutenberg–Richter extrapolation, synthetic validation, and wide-box percentile checks — never from a classifier trained on one positive.

Fault model and stress sources. The active-fault geometry is taken from the GEM Global Active Faults Database (Styron & Pagani, 2020), reduced to 97 segments in the region, each with a strike used both to orient anisotropic aftershock kernels and to define Coulomb receiver planes. Time-causal Coulomb stress changes (ΔCFS) are computed with a pure-numpy Okada (1992) half-space dislocation solution and King, Stein, and Lin (1994) resolution of the stress tensor onto per-segment receivers of the form $\Delta CFS = \Delta\tau + \mu'\Delta\sigma_n$, with rupture dimensions scaled from magnitude via Wells and Coppersmith (1994). Only ruptures preceding the forecast time contribute, preserving causality. Source parameters for the largest ruptures cite the primary studies (Table 2 lists the principal stress sources): Barka et al. (2002) for the 1999 İzmit slip model, Karabulut et al. (2021) for the 2019 Silivri/ Kumburgaz-basin sequence, and Ertuncay et al. (2025) together with the USGS moment tensor for the 2025 Mw 6.2. Interseismic strain is taken as a static field derived from Nevada Geodetic Laboratory GNSS velocities (Blewitt et al., 2018); we treat it as time-invariant by design and later test — and reject — its marginal predictive value. The historical record of M~7 Marmara earthquakes (1509 Mw~7.2 central Marmara; 1719 and 1999 on İzmit; 1754 and 1766 on the Çınarcık/Princes Islands segments; the second 1766 event on Ganos; 1894 on Princes Islands/Yalova; 1912 on Ganos/Mürefte) follows Ambraseys (2002) and Parsons (2004) and anchors the Brownian- Passage-Time renewal layer.

Table 2. *Principal Coulomb stress sources and their primary references.*

Source event	Segment	Reference for parameters
1999 Mw 7.4 İzmit	İzmit / eastern NAF	Barka et al. (2002)
2019 Mw 5.8 Silivri	Central Marmara (Kumburgaz basin)	Karabulut et al. (2021)
2025 Mw 6.2 Kumburgaz	Central Marmara	Ertuncay et al. (2025); USGS moment tensor
Historical M~7 (1509–1912)	all four Marmara segments	Ambraseys (2002); Parsons (2004)
Rupture dimension scaling	—	Wells & Coppersmith (1994)
Half-space dislocation / receiver stress	—	Okada (1992); King et al. (1994)

3. Methods

Features and the leakage self-test. Each (cell, window) carries 19 strictly causal features (Table 1 lists the groups): multi-scale event counts (30/90/365 d, in-cell and in 3- and 5-cell neighbourhoods), times since the last M≥3.5 and M≥4.5 within 25 km, local b-positive, a short-/long-term rate ratio, mean depth, distance to the nearest fault, and the time-causal ΔCFS. All features use only events with time < t0. The central methodological safeguard is a **truncated-catalogue leakage self-test**: for a sample of (cell, t0) rows, every feature is recomputed from a catalogue truncated to < t0 and must reproduce the stored grid to within 0.0 (bit-for-bit), and no feature may correlate > 0.999 with any target . Unlike a target-permutation check, which can miss look-ahead that is not target- correlated, this self-test makes causality a gate any reader can run directly.

Table 1. *The 19 causal features (grouped), each with its rationale and source.*

Group	Features	Rationale / source
Local rate	counts 30/90/365 d in-cell	Omori–Utsu clustering (Ogata, 1988)
Neighbourhood rate	3- and 5-cell counts 30/365 d	spatial smoothing of triggering
Recency	days since last M≥3.5 / M≥4.5 within 25 km	sequence age
Size distribution	local b-positive (50 km, 730 d)	transiently-robust b (van der Elst, 2021)
Rate change	short/long rate ratio	quiescence / activation
Depth	mean depth 90 d	crustal vs deeper activity
Fault geometry	distance to nearest fault	proximity to mapped structure (Styron & Pagani, 2020)
Static stress	time-causal ΔCFS, per-segment receivers	Coulomb triggering (Okada, 1992; King et al., 1994)
ETAS coupling	ln(λ) from the cascade (hybrid grid only)	physics prior for the ML stage

ETAS and a null incompleteness correction. We fit an in-house space–time ETAS by maximum likelihood (Ogata, 1988). The conditional intensity at location **x** and time *t* is λ(**x**, *t*) = μ(**x**) + Σ_{i: t_i < t} K e^{α(m_i − m_c)} (t − t_i + c)^{−p} f(**x** − **x**_i; m_i), a background rate μ plus self-exciting Omori–Utsu aftershock terms with productivity K e^{α(m_i−m_c)}, temporal decay governed by (c, p), and a spatial kernel f whose scale grows with parent magnitude. The background is estimated in the spirit of the stochastic-declustering approach of Zhuang et al. (2002), and the branching structure is validated against Veen and Schoenberg (2008)-style recovery tests. The implementation passes 5/5 unit tests: recovery of known parameters from simulated catalogues, causal filtering (events at or after t0 never affect a forecast), Omori decay of the triggered rate, and Gutenberg–Richter consistency of the simulator's magnitude output. The branching ratio (the expected number of direct offspring per event) is kept subcritical for stable simulation. We attempted a short-term aftershock-incompleteness (STAI) correction using a time-varying completeness Mc(t) (Helmstetter et al., 2006) and raised the branching cap from 0.95 to 0.999 to test whether the maximum-likelihood estimate would leave the cap. **Both attempts returned null results:** at base Mc = 3.0 the post-M5.5 incompleteness window is sub-hour, so the STAI correction dropped 0 events (consistent with the ETASI formulation of Hainzl et al. (2024), which un-pins only at a low base Mc); and raising the cap did not un-pin the fit — the estimate went straight to the new cap (the documented degenerate near-critical mode). We therefore retain the 0.95 cap for stable simulation, applied identically to the feature ETAS and the ETAS baseline .

Cascade Monte-Carlo forecaster. For each window [t0, t0+H) we forward-simulate new background events plus residual-Omori offspring of all prior history plus full recursive cascades, vectorized across K sims (K = 500 in backtests, 10,000 live). Per-cell rare-magnitude rates are computed analytically from the dense M≥3.5 rate field via Gutenberg–Richter, λ_cell(M≥X) = λ_cell(M≥3.5)·10^(−b(X−3.5)), which removes the Monte-Carlo λ6 = 0 artifact. Offspring of M≥5.5 parents are placed with an elliptical 3:1 kernel oriented along the parent's first-72 h aftershock principal axis (or the nearest fault strike); smaller events are isotropic. This design follows the simulation-based operational-forecasting literature (Hardebeck et al., 2024; Mizrahi et al., 2024). The forecaster passes a four-part gate: (a) future events change nothing (causal); (b) the reliability slope on synthetic catalogues is 0.949, within [0.8, 1.2] ; (c) the day-after-M6 rate is 1.79× the first-generation `expected_counts`, quantifying how first-generation ETAS under-counts active sequences; and (d) anisotropy elongates M≥5.5 offspring along-strike 2.77× (vs 1.06 isotropic).

Hybrid and baselines. The hybrid rate is the log-linear blend λ = λ_sim^(1−w)·λ_ML^w, where λ_sim is the cascade rate, λ_ML is a Poisson gradient-boosted regression on the 19 features plus ln(λ_sim), and the single scalar weight w ∈ [0, 1] is selected on the **validation** per-event log-likelihood. The blend is conservative: because w is bounded and chosen out-of-sample, the hybrid degenerates to the pure cascade when the machine-learning stage adds no validated value (w → 0), so it *structurally cannot lose* to its physical core on the validation objective — the ML stage can only correct the cascade where a correction generalises. Monotonicity constraints tie higher ETAS rate and shorter fault distance to higher predicted probability, preventing the booster from learning non-physical inversions. We compare against four baselines — Poisson climatology, fault-proximity, smoothed seismicity, and the ETAS simulation — and, separately, first-generation ETAS `expected_counts`, so that the cascade's contribution is isolable from the ML stage's. Windows are

split by start time into training (≤ 2021-12-31), validation (2022–2023), and test (2024-01-01 onward, with the target window fully inside reviewed data, to + 30 d ≤ 2026-03-31); the split is strictly temporal so that no future window informs a past one. The preliminary catalogue tail is used only by the live forecast, never in any metric.

Metrics. The headline metric is information gain (IG) in nats per observed event relative to a baseline — the mean difference in Poisson log-likelihood per event between the model and the baseline, IG = (1/N) Σ [ln L_model − ln L_base] — computed through one shared scoring function for the model and every baseline so that no predictor enjoys a private evaluation path (Rhoades et al., 2011). A positive IG means the model assigns more probability, per observed event, to where events actually occurred. We also report precision–recall AUC (the discriminative metric most sensitive to the rare positive class), ROC-AUC, the Brier score, and Molchan/area-skill statistics (Zechar & Jordan, 2008), plus reliability diagrams for calibration. Because IG and PR-AUC emphasise different things — likelihood mass versus rank ordering — we treat a result as robust only when the two agree, and flag cases (such as the sparse model-box y45) where they do not.

Synthetic large-event layer. Because the model box contains a single M6, we train a conditional "bigger-event-ahead" discriminator on synthetic sequences. From 250 base ETAS simulations, each seeded with Brownian-Passage-Time-timed (Matthews et al., 2002) characteristic mainshocks of M ~ U(6.8, 7.6) on the four segments (each with a full aftershock cascade), we take 31,712 snapshots 1/3/7 days after every simulated M≥4.5; the label is a subsequent simulated event ≥ (largest-so-far − 0.3) within 30 days. Features are seismicity-only (b-positive drop, sequence count, largest magnitude, time since start, distance to nearest segment) .

Alert and renewal layers. A Foreshock Traffic-Light System (Gulia & Wiemer, 2019) flags relative b-value drops within a sequence; we treat it as an alert layer, not a rate multiplier, and print the low-coverage pseudo-prospective caveat established for the FTLS (Gulia et al., 2020). A Brownian-Passage-Time renewal layer (Matthews et al., 2002; Parsons, 2004) supplies segment-level 30-year conditional M~7 probabilities; it is a large-event layer only and is never blended into the y35/y45 products.

Information-arrival analysis. To characterise *when* a target event becomes forecastable, we define a general procedure that is independent of any single event. For a target (time T, location), we freeze the system at a schedule of strictly causal lead times (e.g. T−365 d, T−180 d, …, T−1 h, and just after any foreshock), recompute all products using only data available at each freeze, and track three quantities as functions of lead time: the target cell's spatial percentile among all cells, the per-cell probability *gain* over a uniform baseline, and the conditional discriminator score. The resulting curves separate persistent spatial hazard (flat percentile) from transient temporal information (a jump in gain when an informative event arrives). We apply the analysis to the 2025 Marmara events; every freeze is checked to use only pre-freeze data.

Prospective protocol. A monthly job issues a 30-day forecast at the catalogue end, appends it to a sha256-hashed, append-only log **before** the outcome is known, and scores past forecasts once their windows close, giving genuine out-of-sample credibility that a backtest cannot. The hash chain makes the record tamper-evident: a forecast cannot be silently revised after the fact, so the accumulating log is a commitment device rather than a retrospective fit. The catalogue is advanced each month from the

KOERI preliminary monthly feed; because the preliminary tail is untyped, it enters only the features of the live forecast and never the reviewed-data evaluation.

4. Results

We organise the results from the general to the specific: the catalogue and its completeness, then the M≥3.5 and M≥4.5 forecasting comparisons and their calibration, then the conditional large-event layer, then the information-arrival analysis applied to the 2025 Marmara events, and finally the live products. Throughout, the relevant comparison is against a properly-fit ETAS, not a naive baseline, and the single model-box M6 caps how strong any large-event claim can be.

Catalogue and completeness (Table 3). The processed catalogue and its completeness structure are summarized in Table 3; the key operational fact is the single model-box M6, which forces every large-event number to rest on extrapolation and simulation rather than direct classification.

Table 3. *Catalogue summary (model box, 2003–2026).*

Quantity	Value
Model-box events	31,329
Date range	2003-01-04 – 2026-07-06
Completeness Mc (max-curvature)	3.65
ML-population completeness (`mag_w`)	~2.72
Events M≥3.5 / ≥4.5 / ≥5.5 / ≥6.0	5,021 / 192 / 12 / 1
b-value ensemble (Aki / b-positive / operational)	1.02 / 1.54 / 1.20

M≥3.5 forecasting (Table 4). On the M≥3.5/30-day target (167 test positives), the cascade Monte-Carlo forecaster ranks best (PR-AUC 0.130), narrowly ahead of first-generation ETAS (0.126) and the ETAS×ML hybrid (0.117, w = 0.7) . The information-gain picture summarises the comparison: the hybrid clears every naive baseline by a wide margin (IG vs Poisson +1.14, vs fault-proximity +0.95, vs smoothed +0.51 nats/event) and edges the cascade in likelihood (+0.25), but it **does not beat a properly-fit ETAS** (IG −0.38), and first-generation ETAS retains the best ROC-AUC (0.894) . A plain machine-learning classifier (without the cascade prior) tells the same story from the other side: it clears the naive baselines (IG vs Poisson +1.07) but loses to ETAS (IG −0.46; ETAS PR-AUC 0.125 ≥ classifier 0.101) . These results are consistent with the EarthquakeNPP finding that flexible models do not systematically beat ETAS (Stockman et al., 2025). The value the cascade adds is concentrated *inside active sequences*, where it produces ~1.8× more offspring than first-generation ETAS — the regime where a Marmara sequence would matter.

Table 4. *M≥3.5/30-day, test set (167 positives): ranking and information gain.*

Predictor	PR-AUC	ROC-AUC	Brier	IG vs predictor (hybrid − ·)
Cascade (ETAS-sim)	**0.130**	0.877	0.00498	+0.25
First-generation ETAS	0.126	**0.894**	0.00497	−0.38
Hybrid (cascade × ML, w = 0.7)	0.117	0.884	**0.00495**	—
Smoothed seismicity	0.044	0.886	0.00544	+0.51
Poisson climatology	0.032	0.848	0.00532	+1.14
Fault-proximity	0.010	0.687	0.00526	+0.95

M≥4.5 overfit and its fix. On the model box (only 22 test positives) the hybrid overfits: PR-AUC 0.033 < cascade 0.066 < first-generation ETAS 0.067, and the hybrid Brier (0.00121) is *worse* than the cascade's (0.00068), so the large IG-versus-cascade there (+1.30) is an artifact . Training instead on the wide box (593,775 rows, 201 M≥4.5 positives) and evaluating on the model box only drives the selected weight down to w = 0.1 (leaning on the cascade), improves the Brier to 0.00078, and yields IG(hybrid vs cascade) = +0.735 with PR-AUC 0.022 vs 0.008 — information gain and PR-AUC now agree . The production y45 product is the wide-box hybrid; the model-box y45 is diagnostic-only.

Calibration (Table 5). A pseudo-prospective calibration battery (2022 → 2026-03) finds every product calibrated in the mean: obs/exp is 1.08 for y35, 1.06 for y45, 0.98 for P(M≥5), 1.16 for P(M≥5.5), and 1.43 for P(M≥6) — the last being a single observed event against 0.70 expected, within Poisson noise . In shape, y35 is calibrated (reliability slope 1.067); the model-box y45 is marginally overconfident (slope 0.785, isotonic on 13 validation positives being insufficient), which is exactly why the wide-box w = 0.1 variant is the production choice; the M≥5/5.5/6 products are too rare per cell to bin and are reported through the aggregate only. No recalibration erodes skill: at y35 the cascade and hybrid keep positive IG against Poisson, fault-proximity, and smoothed baselines.

Table 5. *Calibration battery, pseudo-prospective 2022 → 2026-03 (mean obs/exp).*

Product	Observed	Expected	Obs/Exp	Reliability slope
y35 (M≥3.5, 30 d)	277	256.3	1.08	1.067 (calibrated)
y45 (M≥4.5, 30 d)	35	33.1	1.06	0.785 (→ wide-box variant)
P(M≥5.0, 30 d)	10	10.19	0.98	too rare to bin
P(M≥5.5, 30 d)	3	2.59	1.16	too rare to bin
P(M≥6.0, 30 d)	1	0.70	1.43	too rare to bin

Conditional large-event discriminator (Figure 2). The synthetic discriminator reaches test PR-AUC 0.180 and ROC-AUC 0.738 — about 3.6× the base rate — and is reasonably reliable across its output range (Figure 2) . Applied to the real 2025 sequences with no retraining, it ranks the escalating sequence far higher: the escalating Sındırgı doublet scores 0.277 (33rd percentile) versus the decaying Kumburgaz sequence at ~0.01 (1st–2nd percentile) . Consistent with the physics, the same synthetic framework returns a clear **null** for short-term M≥5.5 predictability from seismicity features (PR-AUC ~0.03, IG ≈ 0 vs sim-Poisson): even with ETAS ground truth, large-event timing is close to a clustering-modulated Poisson process.

Information-arrival analysis of the 2025 events (Figures 3–4). We apply the information-arrival method to three 2025 Marmara events, which between them exercise the method's diagnostic range. The 23 April Mw 6.2 is the primary example (Figure 3). Its spatial part is solved: at a strictly causal freeze the day before, the epicentral cell was the 99.1th-percentile cell by P(M≥6), a gain of 4.55× over a uniform baseline and in the top decile (Figure 4) . Its temporal part shows near-unforecastability: across the freeze schedule (T−365 d, T−180 d, …, T−1 h) the epicentral percentile stays flat at ~98–99 all year (persistent fault hazard) while the 30-day M≥6 probability gain stays ~4–7× — and the cell was quiescent throughout (0–2 M≥2 events per 30 d). Only the lone M4.5 foreshock changed the picture: ten minutes after it, the cell jumped to the 99.9th percentile and the M≥6 probability gain rose to 42× (30-day P(M≥6) from ~0.007% to ~0.060%), with the discriminator rising to 0.76 . The same analysis applied to the neighbouring 2 October 2025 M5.0 shows a different regime: there the discriminator did respond earlier (0.75–0.90 in the weeks before), consistent with a more sequence-preceded event, though it remains noisy. The 10 August 2025 Sındırgı M6.1 exercises the method's coverage limit: it sits **south of the model box** and is therefore not spatially covered (percentile 0 — a coverage boundary, not a forecasting failure), and its out-of-box discriminator score was a weak 0.09–0.28 . Across the three, the method returns the same structural reading — *where* is solved; *when* is bounded by foreshock physics — and quantifies that boundary event by event.

Live products and horizons. At t0 = 2026-07-05, the live 30-day regional P(M≥6) is 0.13%–4.56% across the b-ensemble, with a central (b_op) value of 1.12% (the spatial M≥3.5 product is shown in Figure 1) — a fraction of a percent per month, the same order as a plain-Poisson base rate. The highest 30-day M≥3.5 cell is Marmara Denizi (28.35° E, 40.85° N, P ≈ 43%), the 2025 aftershock zone. Quarter- and year-ahead products are backtest-scaled: four annual backtests give predicted 51.4 vs realized 62 M≥3.5 counts (reliability slope 1.21, within [0.7, 1.3]), so the global scaling factor is 1.0 ; the central-b_op P(M≥6) is 3.9% over a quarter and 15.7% over a year.

Renewal and extension sources. The Brownian-Passage-Time renewal layer gives 30-year conditional M~7 probabilities of 0.1% for İzmit (ruptured 1999), 8.1% for Ganos, 13.7% for Princes Islands, and 21.8% for Central Marmara/Kumburgaz; the combined Princes + Central 30-year probability of 32.5% is within the Parsons (2004) 30–50% ballpark (sanity anchor passes) . A source-extension harness returns a null result for static GNSS strain coupling: a leakage-clean `gnss_rate_change` feature gives IG of −0.014 (validation) and −0.063 (test) — no measurable gain, as expected from a time-invariant field .

5. Discussion

The clearest result of this study is a **where-vs-when decomposition**. The spatial forecasting problem for Marmara is, in an operational sense, largely solved: the seismicity-feature model concentrates probability on the correct fault cells, and in the Mw 6.2 case the epicentral cell was persistently in the top ~1% — an 8× localization relative to a uniform baseline — for a full year before the event. The temporal problem is different in kind. The information-arrival analysis shows that the specific *timing* of the Mw 6.2 was near-unforecastable from the catalogue: the cell was quiescent, and the only short-term information arrived with the lone M4.5 foreshock, ten minutes before the mainshock. The method quantifies this boundary: the probability gain is flat at 4–7× for a year and

then jumps to 42× when the foreshock appears. Run across the three 2025 events, it locates each on the same where-solved / when-foreshock-bounded axis.

This decomposition reconciles our results with the broader literature. The EarthquakeNPP benchmark found that flexible models do not beat ETAS (Stockman et al., 2025); our hybrid, built specifically so it *cannot* lose to its physical core, nonetheless does not beat a properly-fit ETAS in information gain. That is not a failure of engineering but a property of the problem: once leakage is removed, the seismicity-only signal is close to what ETAS already extracts. Where our cascade does add value — inside active sequences, and in the analytic per-cell rare-rate field that supports the countdown — it does so through better physics, not through a more flexible regressor.

A second theme is **conditional forecastability**. The synthetic discriminator ranks sequence-preceded escalation (Sındırgı) far above cold-start decay (Kumburgaz), suggesting that the regime in which short-term large-event forecasting has traction is the sequence-preceded one. This is the same regime the Foreshock Traffic-Light System targets (Gulia & Wiemer, 2019; Gulia et al., 2020), and our FTLS overlay reached red after the Mw 6.2 as expected. The operational implication for Marmara is a **sequence mode**: a persistent spatial hazard map, augmented by sequence-triggered alerts when foreshock-like b-value drops appear — with the explicit understanding that cold-start mainshocks (of which the Mw 6.2 was effectively one) will not be caught in time.

Third, our audit-first design speaks directly to the reliability of multi-source "accuracy" claims in the machine-learning seismicity literature. Studies that evaluate on random train/test splits, or that add exotic data sources without a causality gate, routinely report skill that does not survive forward validation. Our benchmark makes the opposite bet: a smaller, sometimes negative, but *machine-checkably causal* set of numbers. The consequence is that the machine-learning contribution is modest and the physical models are strong.

A fourth point concerns **how the audit changes which results are believable**. Removing leakage changes which numbers are believable, not only their size. The sparse model-box y45 target illustrates this: an information gain of +1.30 against the cascade coexists with a *worse* Brier score and a lower PR-AUC, a combination that is only interpretable as overfitting to 22 positives. A pipeline without a calibration battery and without the discipline of requiring IG and PR-AUC to agree would have reported the +1.30 as skill. The wide-box remedy regularises the weight to w = 0.1 and brings the two metrics into agreement. We report such disagreements — and null results, like the STAI correction and the GNSS feature — as part of the benchmark.

Finally, on **responsible communication**: we report probability *gain* relative to base rates. The live regional 30-day M≥6 probability is a fraction of a percent, consistent with the long-term rate; the renewal layer's decade- scale probabilities for the locked central segment are elevated but consistent with prior hazard estimates (Parsons, 2004). Nothing in these results implies an imminent large earthquake, and the eastward-migration context (Martínez-Garzón et al., 2025) should be communicated as a long-term hazard, not a short-term forecast. The appropriate operational posture is therefore a standing, well-calibrated spatial hazard map that is transparent about its temporal limits, plus a sequence-triggered alert layer whose false-alarm and miss characteristics are stated up front.

6. Limitations and future work

Several limitations bound the results. The live forecast uses a *preliminary* catalogue tail (features only; excluded from every metric), whose magnitudes are untyped. The b-value uncertainty dominates all M6 extrapolation: the mixed Md/ML catalogue inflates b-positive to ~1.54, and the M6 rate can differ several-fold across the ensemble, so we report the ensemble range rather than a point estimate. The study covers a single region with a single real M6, so the large-event results rest on synthetic validation and wide-box percentile checks rather than direct classification. The completeness structure is a genuine mixture (Md and ML populations complete at different magnitudes), which we handle by using a consistent conservative Mc but which limits how low the target magnitude can go. The ETAS fit is isotropic and pinned at the 0.95 branching cap (the STAI and un-pinning attempts returned nulls), and the backtest cascade is limited to K = 500 sims, which bounds the resolution of the rarest per-cell rates. The Sındırgı sequence exposed a hard spatial-coverage limit of the model box.

Future work follows from the null results. The single highest-value data lever is a dense, machine-learning-repicked micro-catalogue (Mc ~ 1.5), which would let the features track nucleation below the current target; time-variable GNSS coupling (rather than a static strain field) and repeating-earthquake creep observables are the next two physical levers, both currently limited by data availability. Finally, the prospective, hashed forecast log should be allowed to accumulate: only an out-of-sample record — which a backtest cannot substitute for — can establish whether the spatial skill demonstrated here translates into operational value.

Data and code availability

All code, the processed KOERI-derived catalogue, the fault and strain inputs, and every results artifact cited above are available in the public repository at https://github.com/keremalhan/marmara-forecast, which reproduces the full pipeline and its four correctness gates from a fresh environment. The catalogue is redistributed as a derived, processed dataset with attribution to KOERI (Boğaziçi University Regional Earthquake-Tsunami Monitoring Center); GNSS velocities are from the Nevada Geodetic Laboratory (Blewitt et al., 2018) and the fault model from the GEM Global Active Faults Database (Styron & Pagani, 2020). See the repository's `DATA_LICENSE.md` for full terms.

Author contributions

Basri Kerem Alhan: Conceptualization, Methodology, Data curation, Writing. **Kenessary Khabat:** Software, Formal analysis, Writing.

Acknowledgements

We thank the Boğaziçi University Kandilli Observatory and Earthquake Research Institute Regional Earthquake-Tsunami Monitoring Center (KOERI-RETMC) for the earthquake bulletin that underlies the catalogue, the Nevada Geodetic Laboratory for the GNSS velocity solutions, and the GEM Global Active Faults Database for the fault model. This research received no external funding.

Figures

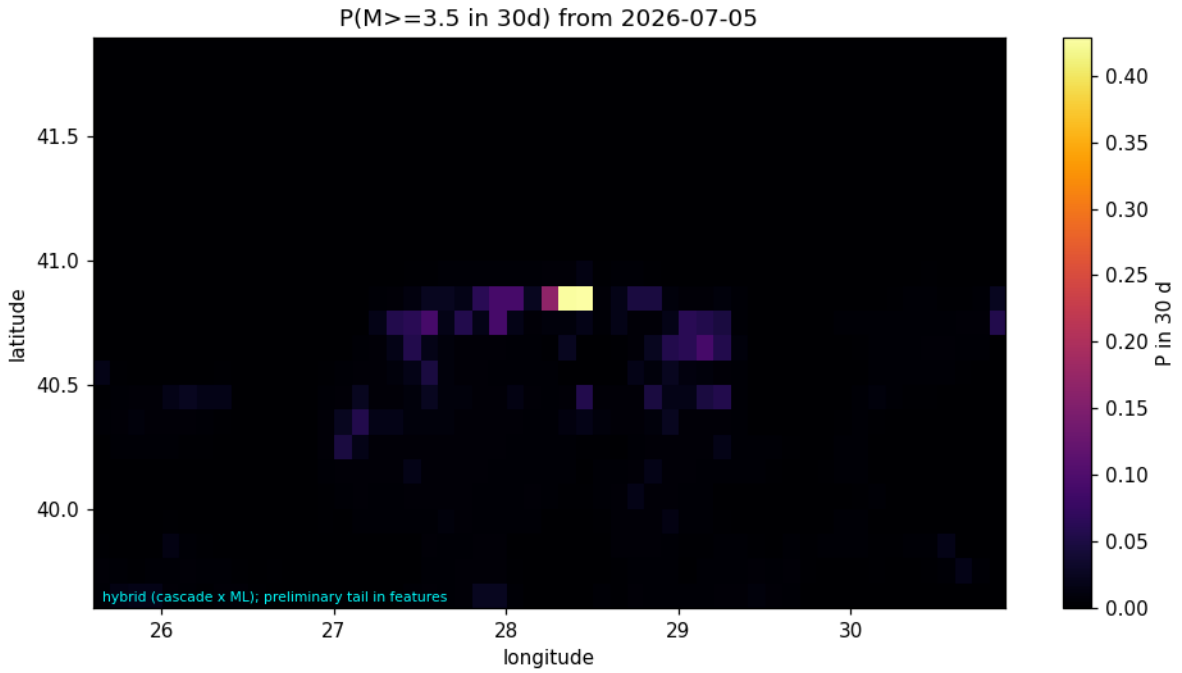

Figure 1. Live 30-day probability of an M≥3.5 event per 0.1° cell at t0 = 2026-07-05. Probability concentrates on the central Marmara segment (the 2025 Mw 6.2 aftershock zone).

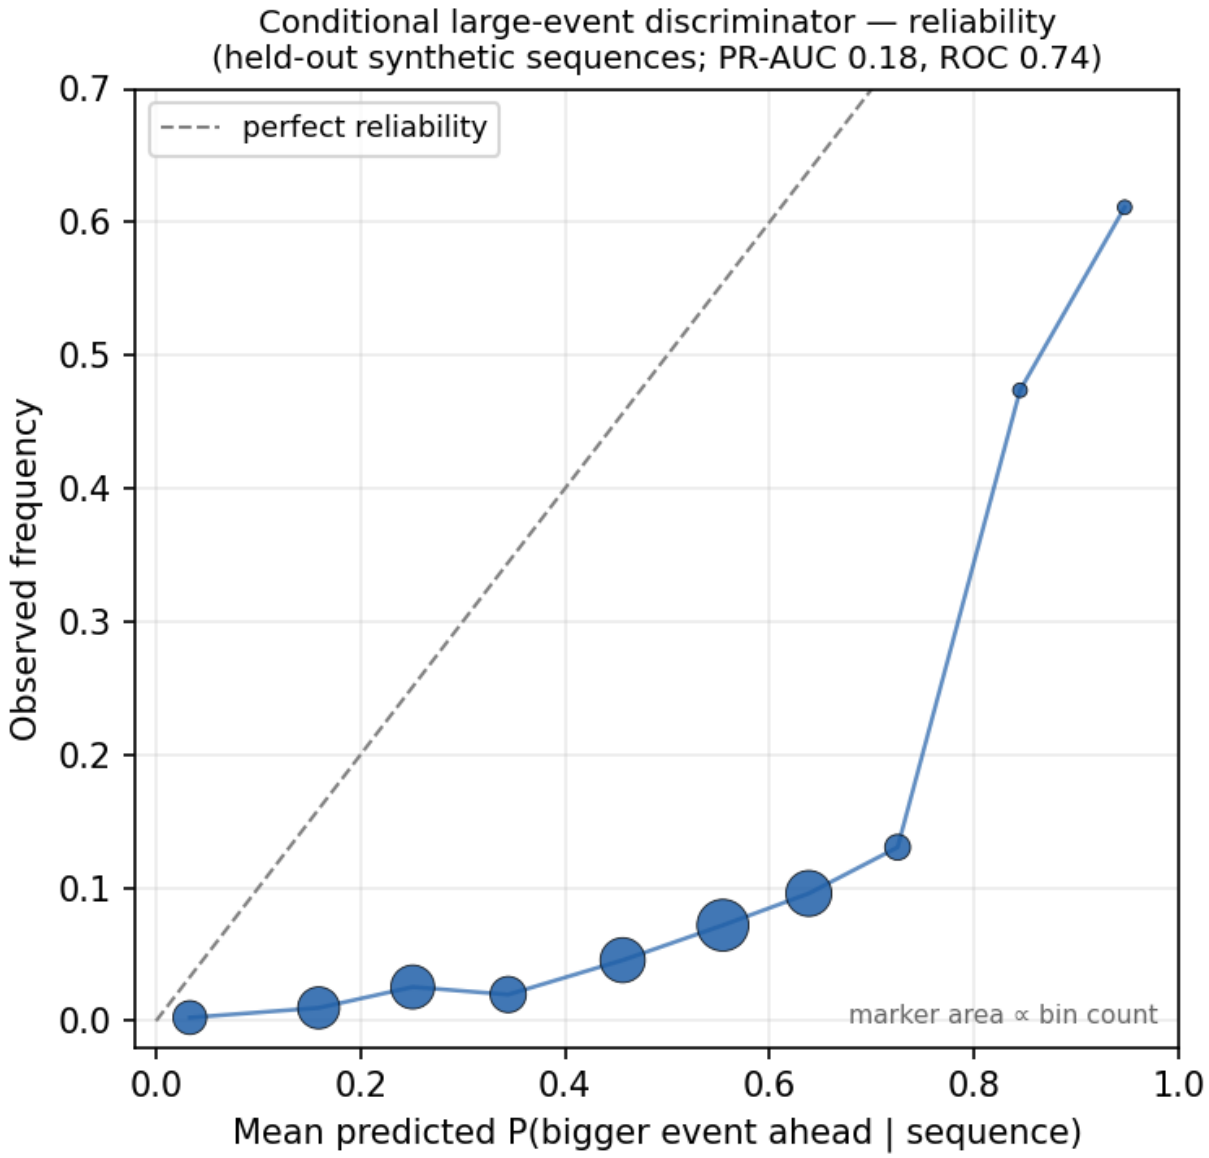

Figure 2. Reliability of the conditional large-event ("bigger-event-ahead") discriminator on held-out synthetic sequences (test PR-AUC 0.18, ROC-AUC 0.74). Marker area is proportional to bin count; the dashed line is perfect reliability.

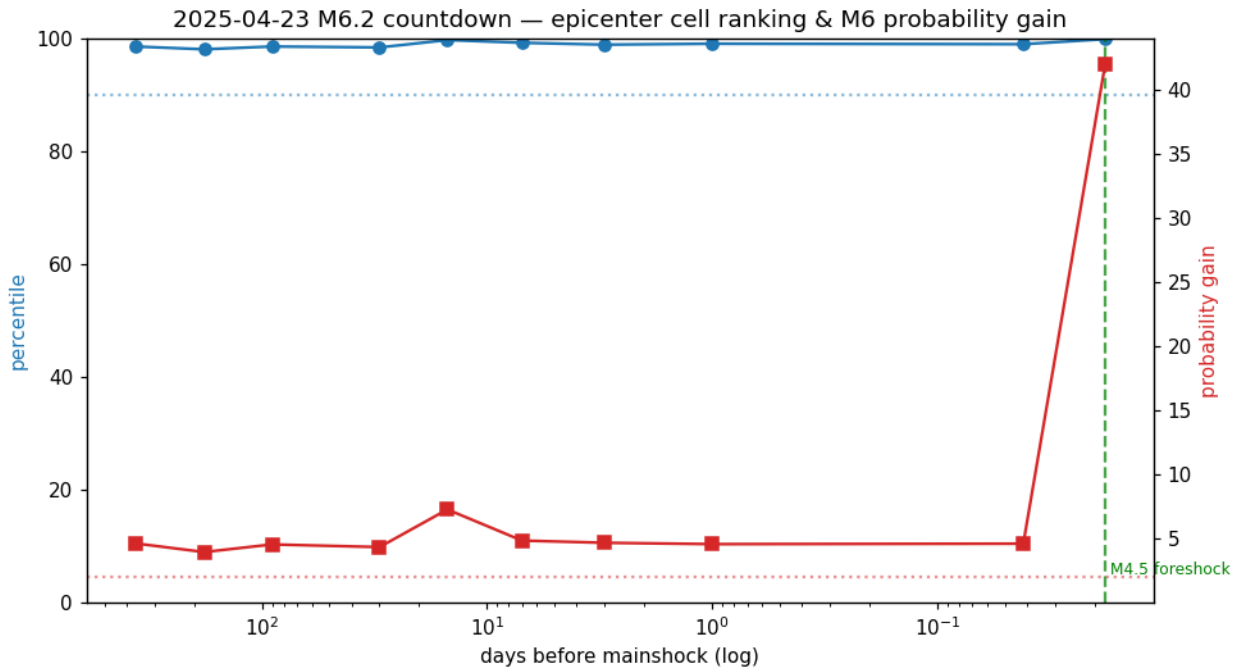

Figure 3. Information-arrival analysis applied to the 23 April 2025 Mw 6.2. Across strictly causal freezes from T−365 d to +10 min, the epicentral-cell percentile stays flat at ~98–99 all year while the 30-day M≥6 probability gain holds at ~4–7×, then jumps to 42× only after the lone M4.5 foreshock, ten minutes before the mainshock. The cell was quiescent throughout. The flat-then-jump shape is the method's signature separation of persistent spatial hazard from transient temporal information.

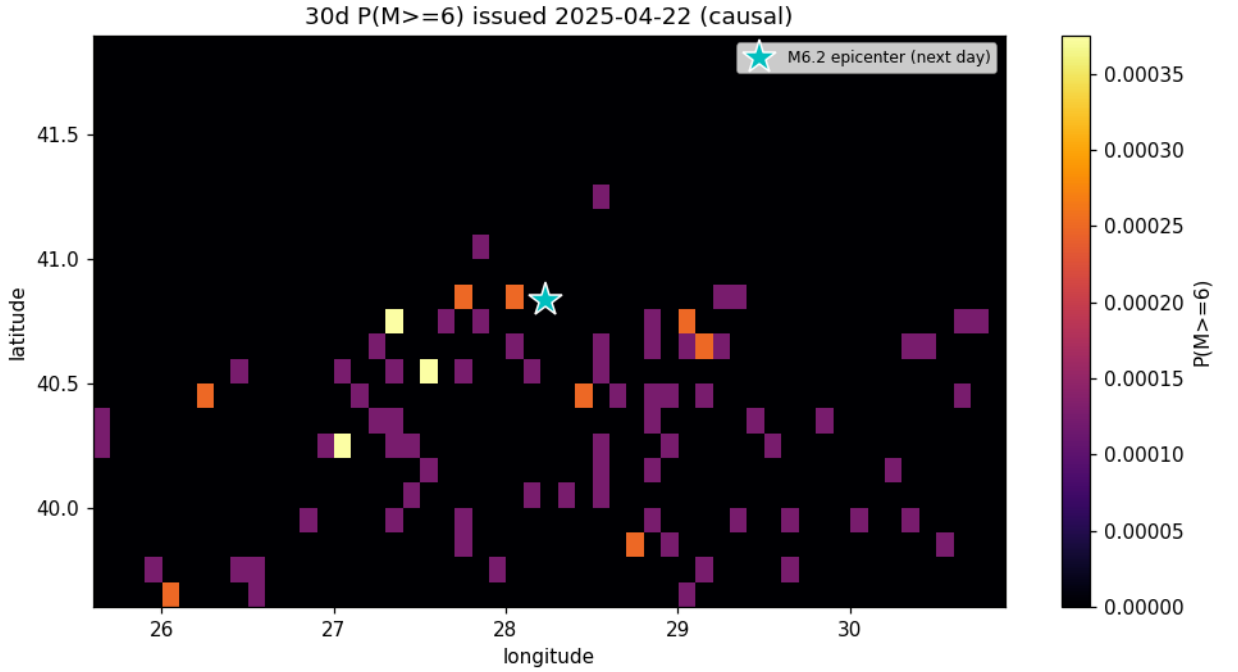

Figure 4. Strictly causal pre-event forecast map for 22 April 2025 (the day before the Mw 6.2). The epicentral cell is the 99.1th-percentile cell by P(M≥6), a 4.55× gain over uniform — the spatial "where" that the information-arrival analysis finds is available far in advance.

References

Aki, K. (1965). Maximum likelihood estimate of b in the formula $\log N = a - bM$ and its confidence limits. *Bulletin of the Earthquake Research Institute, 43*, 237–239.

Ambraseys, N. (2002). The seismic activity of the Marmara Sea region over the last 2000 years. *Bulletin of the Seismological Society of America, 92*(1), 1–18. https://doi.org/10.1785/0120000843

Barka, A., Akyüz, H. S., Altunel, E., Sunal, G., Çakir, Z., Dikbas, A., Yerli, B., Armijo, R., Meyer, B., de Chabalier, J. B., Rockwell, T., Dolan, J. R., Hartleb, R., Dawson, T., Christofferson, S., Tucker, A., Fumal, T., Langridge, R., Stenner, H., … Page, W. (2002). The surface rupture and slip distribution of the 17 August 1999 İzmit earthquake (M 7.4), North Anatolian Fault. *Bulletin of the Seismological Society of America, 92*(1), 43–60. https://doi.org/10.1785/0120000841

Blewitt, G., Hammond, W. C., & Kreemer, C. (2018). Harnessing the GPS data explosion for interdisciplinary science. *Eos, 99*. https://doi.org/10.1029/2018EO104623

Dascher-Cousineau, K., Shchur, O., Brodsky, E. E., & Günnemann, S. (2023). Using deep learning for flexible and scalable earthquake forecasting. *Geophysical Research Letters, 50*(17), e2023GL103909. https://doi.org/10.1029/2023GL103909

Ertuncay, D., et al. (2025). 23 April 2025 Marmara Sea (Mw 6.3), Türkiye earthquake: Mainshock, aftershock, and ground observations. *Seismica*. https://seismica.library.mcgill.ca/article/view/1773

Gulia, L., & Wiemer, S. (2019). Real-time discrimination of earthquake foreshocks and aftershocks. *Nature, 574*(7777), 193–199. https://doi.org/10.1038/s41586-019-1606-4

Gulia, L., Wiemer, S., & Vannucci, G. (2020). Pseudoprospective evaluation of the foreshock traffic-light system in Ridgecrest and implications for aftershock hazard assessment. *Seismological Research Letters, 91*(5), 2828–2842. https://doi.org/10.1785/0220190307

Hainzl, S., Kumazawa, T., & Ogata, Y. (2024). Aftershock forecasts based on incomplete earthquake catalogues: ETASI model application to the 2023 SE Türkiye earthquake sequence. *Geophysical Journal International, 236*(3), 1609–1620. https://doi.org/10.1093/gji/ggae006

Hardebeck, J. L., Llenos, A. L., Michael, A. J., Page, M. T., Schneider, M., & van der Elst, N. J. (2024). Aftershock forecasting. *Annual Review of Earth and Planetary Sciences, 52*, 61–84. https://doi.org/10.1146/annurev-earth-040522-102129

Helmstetter, A., Kagan, Y. Y., & Jackson, D. D. (2006). Comparison of short-term and time-independent earthquake forecast models for southern California. *Bulletin of the Seismological Society of America, 96*(1), 90–106. https://doi.org/10.1785/0120050067

Kadirioğlu, F. T., & Kartal, R. F. (2016). The new empirical magnitude conversion relations using an improved earthquake catalogue for Turkey and its near vicinity (1900–2012). *Turkish Journal of Earth Sciences, 25*(4), 300–310. https://doi.org/10.3906/yer-1511-7

Karabulut, H., Güvercin, S. E., Hollingsworth, J., & Konca, A. Ö. (2021). The moderate size 2019 September Mw 5.8 Silivri earthquake unveils the complexity of the Main Marmara Fault shear zone. *Geophysical Journal International, 224*(1), 377–388. https://doi.org/10.1093/gji/ggaa469

King, G. C. P., Stein, R. S., & Lin, J. (1994). Static stress changes and the triggering of earthquakes. *Bulletin of the Seismological Society of America, 84*(3), 935–953. https://doi.org/10.1785/BSSA0840030935

Martínez-Garzón, P., Bohnhoff, M., Mencin, D., Kwiatek, G., Dresen, G., Hodgkinson, K., Nurlu, M., Kadirioğlu, F. T., & Malin, P. E. (2025). Progressive eastward rupture of the Main Marmara Fault towards Istanbul. *Science*. https://doi.org/10.1126/science.adz0072

Matthews, M. V., Ellsworth, W. L., & Reasenberg, P. A. (2002). A Brownian model for recurrent earthquakes. *Bulletin of the Seismological Society of America, 92*(6), 2233–2250. https://doi.org/10.1785/0120010267

Mizrahi, L., Nandan, S., & Wiemer, S. (2021). The effect of declustering on the size distribution of mainshocks. *Seismological Research Letters, 92*(4), 2333–2342. https://doi.org/10.1785/0220200231

Mizrahi, L., Dallo, I., van der Elst, N. J., Christophersen, A., Spassiani, I., Werner, M. J., Iturrieta, P., Bayona, J. A., Iervolino, I., Schneider, M., Marzocchi, W., Gulia, L., Wiemer, S., & Marti, M. (2024). Developing, testing, and communicating earthquake forecasts: Current practices and future directions. *Reviews of Geophysics, 62*(3), e2023RG000823. https://doi.org/10.1029/2023RG000823

Ogata, Y. (1988). Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association, 83*(401), 9–27. https://doi.org/10.1080/01621459.1988.10478560

Okada, Y. (1992). Internal deformation due to shear and tensile faults in a half-space. *Bulletin of the Seismological Society of America, 82*(2), 1018–1040. https://doi.org/10.1785/BSSA0820021018

Parsons, T. (2004). Recalculated probability of M ≥ 7 earthquakes beneath the Sea of Marmara, Turkey. *Journal of Geophysical Research: Solid Earth, 109*(B5), B05304. https://doi.org/10.1029/2003JB002667

Rhoades, D. A., Schorlemmer, D., Gerstenberger, M. C., Christophersen, A., Zechar, J. D., & Imoto, M. (2011). Efficient testing of earthquake forecasting models. *Acta Geophysica, 59*(4), 728–747. https://doi.org/10.2478/s11600-011-0013-5

Stockman, S., Lawson, D. J., & Werner, M. J. (2025). EarthquakeNPP: A benchmark for earthquake forecasting with neural point processes. *Transactions on Machine Learning Research*. https://arxiv.org/abs/2410.08226

Styron, R., & Pagani, M. (2020). The GEM Global Active Faults Database. *Earthquake Spectra, 36*(S1), 160–180. https://doi.org/10.1177/8755293020944182

van der Elst, N. J. (2021). B-positive: A robust estimator of aftershock magnitude distribution in transiently incomplete catalogs. *Journal of Geophysical Research: Solid Earth, 126*(2), e2020JB021027. https://doi.org/10.1029/2020JB021027

Veen, A., & Schoenberg, F. P. (2008). Estimation of space–time branching process models in seismology using an EM-type algorithm. *Journal of the American Statistical Association, 103*(482), 614–624. https://doi.org/10.1198/016214508000000148

Wells, D. L., & Coppersmith, K. J. (1994). New empirical relationships among magnitude, rupture length, rupture width, rupture area, and surface displacement. *Bulletin of the Seismological Society of America, 84*(4), 974–1002. https://doi.org/10.1785/BSSA0840040974

Wiemer, S., & Wyss, M. (2000). Minimum magnitude of completeness in earthquake catalogs: Examples from Alaska, the western United States, and Japan. *Bulletin of the Seismological Society of America, 90*(4), 859–869. https://doi.org/10.1785/0119990114

Zechar, J. D., & Jordan, T. H. (2008). Testing alarm-based earthquake predictions. *Geophysical Journal International, 172*(2), 715–724. https://doi.org/10.1111/j.1365-246X.2007.03676.x

Zhuang, J., Ogata, Y., & Vere-Jones, D. (2002). Stochastic declustering of space–time earthquake occurrences. *Journal of the American Statistical Association, 97*(458), 369–380. https://doi.org/10.1198/016214502760046925