Coral bleaching as state transition: bimodal prevalence and response-definition dependence of thermal stress metrics across Japan's Monitoring Site 1000 network

Hiroki Fukui

Research Institute of Criminal Psychiatry / Sex Offender Medical Center | Dept. Neuropsychiatry, Kyoto University | fukui@somec.org | ORCID: 0009-0008-7122-522X

Abstract

Coral bleaching is conventionally modelled as a continuous response to cumulative thermal stress, yet the distributional structure of site-level bleaching prevalence has rarely been examined. Here we analyse five years (fiscal years 2020–2024) of standardised bleaching surveys from Japan's Monitoring Site 1000 program, encompassing 2,287 observations across 585 survey points at 26 sites in the Ryukyu Archipelago and adjacent waters. We document three principal findings. First, bleaching prevalence is bimodally distributed in all five years: the intermediate range (20–80%) remains stable at 21–27% of observations, while inter-annual variation is driven by redistribution between low and high domains — a pattern confirmed by Hartigan's dip test and beta mixture modelling. Second, the two major bleaching years — 2022 and 2024 — are qualitatively distinct: 2022 was a partial mass bleaching (positively skewed, selective), while 2024 was a comprehensive mass bleaching (symmetric distribution, median 60.0%). Third, the relative performance of thermal stress metrics was strongly response-definition dependent: a simple threshold metric (days above 30°C) exceeded Degree Heating Weeks in discriminating moderate-to-severe bleaching (AUC 0.846 vs 0.758 at ≥50% prevalence; DeLong test, p < 0.001) but not for detection-limit occurrence (>0%), where both metrics performed comparably (AUC 0.748 vs 0.761; p = 0.214). This response-definition dependence indicates that metric structure and the ecological severity threshold defining the outcome are inseparable design considerations — a finding with implications for statistical modelling, event classification, and thermal stress metric development.

---

1. Introduction

Mass coral bleaching, driven by marine heatwaves of increasing frequency and severity, has become the dominant threat to tropical and subtropical reef ecosystems worldwide (Hughes et al. 2018a). The standard predictive framework centres on Degree Heating Weeks (DHW), a cumulative thermal stress metric that integrates the magnitude and duration of sea surface temperatures exceeding a site-specific climatological maximum (Liu et al. 2014). The logic of DHW is inherently cumulative: thermal stress is treated as a dose that accumulates continuously, and bleaching is predicted to intensify as a graded function of that dose. DHW thresholds of 4 and 8 °C-weeks are operationally interpreted as indicating the onset and severe phases of mass bleaching, respectively. But this cumulative logic carries a linearity assumption that is rarely made explicit: it presupposes that bleaching prevalence increases progressively with heat exposure. If instead bleaching is a state transition — a threshold-mediated switch between an unbleached and a bleached state — then cumulative heat dose may be the wrong predictor category altogether.

Despite its operational utility, the DHW framework rests on two assumptions that have received limited empirical scrutiny at the site level. The first is linearity: DHW treats thermal stress as a dose that accumulates continuously, implying that bleaching prevalence should increase as a graded function of cumulative heat exposure. The second is universality: standard thresholds are applied globally, although

regional studies have demonstrated substantial variation in effective thresholds (Sakai et al. 2019; Whitaker & DeCarlo 2024) and in the relative performance of alternative metric formulations (Lachs et al. 2021).

The distributional structure of bleaching prevalence — how site-level values are distributed across a monitoring network — provides a direct test of the linearity assumption. If bleaching is a continuous dose-response, one would expect prevalence values to be unimodally distributed, spreading progressively with increasing thermal stress. If instead bleaching involves a threshold-mediated state transition, as suggested by bioenergetic models of the photodamage–symbiont expulsion cascade (Pfab et al. 2024), one would predict a bimodal distribution: sites cluster into low and high bleaching states with few observations at intermediate values, and inter-annual variation is expressed as redistribution between states rather than as a shift along a continuum.

Testing this prediction requires a dataset with three properties: continuous (not binary) measurement of bleaching prevalence at the site level, standardised protocols across sites, and repeated observations over multiple years encompassing both mild and severe bleaching events. Japan's Monitoring Site 1000 coral reef survey program meets all three criteria. Administered by the Ministry of Environment, the program monitors 26 sites spanning approximately 6° of latitude, from the Yaeyama Islands (24.3°N) to Kushimoto (33.5°N) — the northern distribution limit of hermatypic corals in the northwestern Pacific. Bleaching prevalence is recorded as a continuous percentage at each of 585 survey points, and the five fiscal years 2020–2024 encompass two major bleaching events (2022 and 2024) bracketing three milder years. Despite more than two decades of continuous monitoring since 2003, the Monitoring Site 1000 coral reef dataset has not been the subject of an international peer-reviewed study (Kawagoe 2017 reported 2016 bleaching in the journal of the Japanese Coral Reef Society; Kimura et al. 2022 incorporated the data into an ICRI regional report). The present analysis represents the first such use.

In this study, we examine the distributional structure of bleaching prevalence across this network and report three findings. First, bleaching prevalence is bimodally distributed in all five years, with a stable intermediate zone and inter-annual variation driven by redistribution between low and high domains. Second, the 2022 and 2024 bleaching events are qualitatively distinct — partial versus comprehensive mass bleaching — a distinction invisible to scalar summary statistics. Third, the predictive performance of thermal stress metrics depends on both metric structure and response definition: a simple threshold exceedance metric (days above 30°C) discriminates moderate-to-severe bleaching more effectively than DHW, but not bleaching at the detection limit, where both metrics perform comparably — indicating that metric structure and the ecological severity of the predicted outcome are inseparable design considerations.

Abstract

2. Methods

2.1 Study area and data source

We analysed coral bleaching data from the Monitoring Site 1000 coral reef survey program, administered by the Ministry of Environment of Japan. The program monitors coral reef condition across the Ryukyu Archipelago and adjacent areas using a standardised spot-check protocol (Ministry of Environment 2024). The dataset used in this study covers fiscal years 2020–2024 (approval number 環生多発第2602271号) and comprises 585 survey points across 26 sites, yielding 2,287 observations.

Survey sites span approximately 6° of latitude, from the Yaeyama Islands (24.3°N) to Kushimoto (33.5°N), encompassing the northern distribution limit of hermatypic corals in the northwestern Pacific. Sites are structured as clusters of survey points (typically 10–40 points per site) within a ~5 km² area.

2.2 Bleaching and mortality assessment

At each survey point, bleaching prevalence (bleaching*all) was recorded as the percentage of colony area showing signs of bleaching, averaged across quadrats within the point. Mortality prevalence (mortality*all) was recorded using the same protocol. Both variables yield values at approximately 0.5% resolution (135 unique values across the dataset), reflecting the original recording protocol. Surveys are conducted annually in autumn (September–November), approximately 1–3 months after the summer thermal stress peak. Consequently, the mortality values recorded represent the state at the time of survey and may underestimate total post-bleaching mortality, which can progress over six months or longer (Sakai et al. 2019). This temporal constraint applies to all mortality results reported in this study.

2.3 Sample size variation

The number of survey sites ranged from 21 to 25 per year, with 428–496 survey points per year (Table 1). Variation was driven primarily by the addition or removal of entire sites: three sites (02, 08, 20) were added in FY2022, two (02, 08) were discontinued in FY2023, and site 25 was added in FY2024. Minor variation arose from changes in point numbers within sites (e.g., site 05 ranged from 32 to 42 points). Missing data were negligible (1 record in FY2024). Twenty-one of 26 sites participated in all five fiscal years; sensitivity analyses restricted to this balanced panel are reported in the Supplementary Material (Fig. S1).

Table 1. Sample sizes by fiscal year.

Fiscal year	Sites	Survey points
2020	21	428
2021	22	438
2022	25	496
2023	22	455
2024	23	470
Total	26	2,287

2.4 Satellite sea surface temperature

Daily sea surface temperature (SST) was obtained from the NASA MUR SST v4.1 product (Multi-scale Ultra-high Resolution SST; JPL MUR MEaSUREs Project 2015), which provides global coverage at 0.01° (~1 km) resolution. SST values were extracted for each monitoring site using the nearest grid point to the site centroid.

2.5 Thermal stress metrics

Two thermal stress metrics were computed from the same SST product (MUR SST v4.1), so that the comparison between them reflects the difference in metric structure rather than differences in input data or spatial resolution.

Degree Heating Weeks (DHW) were calculated following the standard NOAA Coral Reef Watch methodology (Liu et al. 2014). The Maximum Monthly Mean (MMM) climatology was calculated as the highest monthly mean SST over the 2003–2014 baseline period — the longest span of complete annual cycles available in MUR at the time of analysis (MUR v4.1 production begins mid-2002). HotSpots were defined as max(SST − MMM, 0), accumulated over an 84-day (12-week) window, and divided by 7 to

yield DHW in °C-weeks. This implementation differs from the operational NOAA CRW algorithm in one respect: the standard CRW algorithm applies a HotSpot > 1°C cutoff, accumulating only anomalies exceeding MMM + 1°C (Liu et al. 2014), whereas our implementation accumulates all positive anomalies above MMM. This choice follows the finding by Lachs et al. (2021) that HotSpot thresholds at or below MMM improve predictive accuracy. An alternative 56-day (8-week) window following Lachs et al. (2021) was also computed; results are reported in the Supplementary Material.

Threshold exceedance days were defined as the number of days with MUR SST exceeding 30°C over the summer period preceding each survey. This threshold was adopted a priori based on Sakai et al. (2019), who identified 30°C as associated with bleaching onset at Sesoko Island, Okinawa.

2.6 Statistical analyses

Distributional analysis

Violin plots of bleaching prevalence and mortality prevalence were generated using Gaussian kernel density estimation with Scott's rule bandwidth selection (scipy.stats.gaussian_kde; Scott 1992). Box plots showing interquartile range and median are overlaid within each violin.

Tests for bimodality

Bimodality of bleaching prevalence was assessed using two independent approaches. First, Hartigan's dip test (Hartigan & Hartigan 1985) was applied to the full data (0–100%) and to the intermediate range (0 < bleaching < 100%) separately, to confirm that bimodality is not an artefact of boundary inflation at 0% and 100%.

Second, one-component and two-component beta mixture models were fitted to the intermediate-range data (0 < bleaching < 100%, scaled to [0, 1]) using an expectation-maximisation (EM) algorithm with three initial parameter sets. Model selection was based on the Bayesian Information Criterion (BIC), with ΔBIC > 10 indicating very strong support for the more complex model (Kass & Raftery 1995).

Intraclass correlation

Intraclass correlation coefficients (ICC) were estimated from a one-way random-effects ANOVA model with site as the grouping factor, using bleaching prevalence in percentage scale. Results were robust to logit-transformed (0, 1) scaling with ϵ-adjustment for boundary values.

Predictive comparison

The predictive accuracy of DHW and threshold exceedance days for bleaching was compared using ROC analysis and the area under the curve (AUC). Both metrics were derived from the same SST product (MUR), isolating the effect of metric structure on predictive performance. Because satellite SST was available for 21 of 26 monitoring sites (five sites lacked MUR coverage), the predictive comparison was conducted on these 21 sites across all five fiscal years (n = 2,142 observations). Three binary response definitions were used: (1) bleaching occurrence (bleaching$all > 0\%)$, (2) moderate bleaching (bleaching$all \geq 50\%$), and (3) severe bleaching (bleaching_all $\geq$ 80%). This multi-threshold approach tests whether the relative performance of predictors depends on the ecological severity of the response, addressing the concern that near-detection-limit observations (0.5%) differ fundamentally from ecologically significant bleaching. Pairwise comparison of AUC values was performed using the DeLong test (DeLong et al. 1988), which accounts for the correlation between ROC curves derived from the same sample. Confidence intervals (95%) were estimated by bootstrap resampling (n = 2,000, seed = 42).

The analyses in this study are descriptive and exploratory in character. The 30°C threshold was adopted a priori from Sakai et al. (2019); all other thresholds (27–32°C) and the alternative 8-week DHW window were examined as robustness checks rather than independent confirmatory tests. No correction for multiple comparisons was applied.

Conditional mortality analysis

Mean and median mortality were computed for five bleaching prevalence bands (0%, 0–20%, 20–50%, 50–80%, ≥80%) to characterise the bleaching–mortality relationship.

Spot-level tracking

Survey points with data in consecutive years were tracked to quantify recovery and cumulative impact trajectories. Specifically, 439 points with data in both 2022 and 2023 were used for recovery analysis, and 435 points with data in both 2022 and 2024 were used for back-to-back bleaching analysis.

All analyses were conducted in Python 3.11 using pandas, scipy, numpy, and matplotlib. Code is available at https://github.com/hirokifukui/marine-obs

---

3. Results

3.1 Bleaching prevalence across five years

Bleaching prevalence varied substantially across the five-year study period (Fig. 1, Table 2). The years 2020, 2021, and 2023 showed low overall bleaching, with median prevalence of 0% and 38–49% of survey points recording any bleaching. In contrast, 2022 and 2024 were major bleaching years, with 78.0% and 85.1% of points recording bleaching, respectively. The year 2024 was the most severe, with a median bleaching prevalence of 60.0% and a mean of 48.4%.

Table 2. Bleaching prevalence (bleaching_all, %) by fiscal year.

Year	n	Median	Mean	Q25	Q75	>0%
2020	428	0.0	17.0	0.0	27.5	49.3%
2021	438	0.0	16.8	0.0	32.1	38.4%
2022	496	20.0	37.6	0.5	80.0	78.0%
2023	455	0.0	12.3	0.0	23.8	44.8%
2024	470	60.0	48.4	2.5	90.0	85.1%

3.2 Bimodal distribution of bleaching prevalence

The violin plots revealed a bimodal distribution in all five years (Fig. 1). Survey points clustered into a low-bleaching domain (<20%) and a high-bleaching domain (≥80%), with the intermediate range (20–80%) remaining stable at 21–27% of observations regardless of year (Table 3). Inter-annual variation was driven by redistribution between the low and high domains: in 2024, the high domain contained 39.4% of observations, compared to 2.7% in 2021.

Table 3. Distribution of survey points across bleaching domains.

Year	Low (<20%)	Intermediate (20–80%)	High (≥80%)
2020	68.9%	27.1%	4.0%
2021	72.4%	24.9%	2.7%
2022	49.0%	21.4%	29.6%
2023	73.4%	25.5%	1.1%
2024	37.4%	23.2%	39.4%

Hartigan's dip test rejected the null hypothesis of unimodality in all years, both for the full data and for the intermediate range (0 < bleaching < 100%; all p < 0.001; Table 4). The dip statistic was largest in 2024 (D = 0.132, full data), consistent with the most pronounced bimodality in the most severe bleaching year.

Table 4. Hartigan's dip test results.

Year	Full data D	Full data p	Intermediate D	Intermediate p
2020	0.044	<0.001	0.055	<0.001
2021	0.037	<0.001	0.097	<0.001
2022	0.091	<0.001	0.116	<0.001
2023	0.055	<0.001	0.063	<0.001
2024	0.132	<0.001	0.105	<0.001

Two-component beta mixture models were favoured over one-component models in every year (ΔBIC = 9.0–113.9; Table 5). Support was very strong (ΔBIC > 10) in four of five years; in 2023, ΔBIC = 9.0 represents strong evidence (Kass & Raftery 1995), but the qualitative pattern — a low-bleaching component and a high-bleaching component — was consistent across all years including 2023. The two components consistently showed a low-bleaching group (α■ < 1, β■ > 2; J-shaped) and a high-bleaching group (α■ > β■ in major bleaching years; right-skewed).

Table 5. Beta mixture model comparison (intermediate range, 0 < bleaching < 100%). ΔBIC > 10 indicates very strong support for the two-component model (Kass & Raftery 1995); 2023 (ΔBIC = 9.0) represents strong evidence.

Year	n	BIC■	BIC■	ΔBIC	w■	α■	β■	w■	α■	β■
2020	211	−82.4	−113.7	31.4	0.740	0.69	3.07	0.260	23.66	7.41
2021	168	−24.0	−79.1	55.2	0.396	0.47	2.94	0.604	6.69	3.56
2022	379	−120.7	−224.3	103.6	0.503	0.75	5.47	0.497	7.14	1.42
2023	204	−148.1	−157.0	9.0	0.701	0.56	2.06	0.299	8.37	10.81
2024	381	−85.3	−199.2	113.9	0.384	0.60	4.74	0.616	6.17	1.49

3.3 Qualitative distinction between the 2022 and 2024 bleaching events

Although both 2022 and 2024 were major bleaching years, their distributional signatures differed (Fig. 1). In 2022, the distribution was positively skewed (skewness = 0.413), with most points in the low-to-moderate range and a tail of severe cases. In 2024, the distribution was approximately symmetric (skewness = −0.100), with the median at 60.0% and the high-bleaching component of the beta mixture model dominant (w■ = 0.616 vs 0.497 in 2022).

3.4 Non-linear relationship between bleaching and mortality

Mortality prevalence was structurally distinct from bleaching prevalence: right-skewed rather than bimodal, with the bulk of observations below 20% even in severe bleaching years (Table 6).

Table 6. ** Mortality prevalence (mortality_all, %) by fiscal year.

Year	n	Median	Mean	>0%
2020	428	0.0	0.9	36.2%
2021	438	0.0	0.6	29.9%
2022	496	2.5	15.2	65.9%
2023	455	0.0	0.2	12.1%
2024	470	10.0	18.7	75.3%

Conditional mortality analysis revealed a non-linear relationship with bleaching prevalence (Table 7). Mortality escalated sharply above 80% bleaching, yet even at 80–100% bleaching, the median mortality was only 25.0%.

Table 7. ** Conditional mortality by bleaching prevalence band (pooled 2020–2024).

Bleaching band	n	Mortality mean	Mortality median	Mortality >0%
0%	917	0.0	0.0	1.1%
0–20%	448	1.5	0.5	50.9%
20–50%	271	4.8	1.0	69.7%
50–80%	285	13.2	2.0	84.9%
≥80%	366	30.9	25.0	96.4%

3.5 Spatial structure of bimodality

The bimodal distribution was driven primarily by between-site divergence rather than within-site variability. In 2024, six sites (11, 12, 18, 19, 24, 25) had all points in the low domain (<20% bleaching), while seven sites (6, 10, 13, 14, 15, 16, 17) had all points in the high domain (≥80%). Intraclass correlation coefficients (ICC) confirmed this spatial partitioning: site identity explained 73–89% of the variance in bleaching prevalence across years (Table 5b). ICC was highest in 2021 (0.886), a mild year in which site-level differences in baseline vulnerability were most apparent, and lowest in 2024 (0.734), when the comprehensive event partially overrode site-level variation.

Table 5b. ** Intraclass correlation coefficients for bleaching prevalence by fiscal year.

Year	ICC	n sites	n obs
2020	0.796	20	411
2021	0.886	20	412
2022	0.841	21	432
2023	0.853	21	439
2024	0.734	21	445

3.6 Recovery and back-to-back bleaching

Bleaching prevalence declined sharply from 2022 (mean 37.6%) to 2023 (mean 12.3%). Among 192 survey points with ≥50% bleaching in 2022, 28.1% recorded zero bleaching in 2023, and 85.4% showed zero mortality at the 2023 survey.

Tracking 435 points across 2022 and 2024 revealed three trajectories: 32.0% experienced ≥50% bleaching in both years; 9.4% transitioned from 0% in 2022 to ≥50% in 2024; 5.5% showed the reverse (≥50% to

0%).

3.7 Predictive comparison: threshold exceedance versus DHW

The predictive comparison was conducted on the 21 sites with satellite SST coverage across all five fiscal years (n = 2,142 observations; see Section 2.6). The predictive performance of thermal stress metrics depended on the response definition (Table 8). For bleaching occurrence at the detection limit (>0%), DHW and days above 30°C performed comparably (AUC 0.761 vs 0.748; DeLong test, z = 1.243, p = 0.214). For moderate bleaching (≥50%), days above 30°C was significantly superior (AUC 0.846 vs 0.758; z = −8.495, p < 0.001). For severe bleaching (≥80%), both metrics performed well with no statistically significant difference (AUC 0.854 vs 0.843; z = −1.038, p = 0.299).

Table 8. **Predictive accuracy of thermal stress metrics across response definitions (21 sites with continuous MUR SST coverage over all five fiscal years, n = 2,142). AUC values with 95% bootstrap confidence intervals; pairwise comparison by DeLong test.**

Response definition	DHW AUC (95% CI)	Days >30°C AUC (95% CI)	DeLong z	p
>0%	0.761 (0.739–0.781)	0.748 (0.726–0.769)	1.243	0.214
≥50%	0.758 (0.735–0.781)	0.846 (0.829–0.862)	−8.495	<0.001
≥80%	0.843 (0.823–0.862)	0.854 (0.836–0.874)	−1.038	0.299

Both metrics were derived from the same SST product, so differences in AUC reflect metric structure rather than input data quality. Notably, DHW's AUC was similar across the >0% and ≥50% response definitions (0.761 and 0.758), while the threshold metric's AUC increased sharply from 0.748 to 0.846 — consistent with threshold exceedance logic capturing a binary state transition that becomes ecologically manifest only above a minimum severity. However, only the difference at ≥50% was statistically significant (p < 0.001); at >0% and ≥80%, the two metrics were not distinguishable in predictive accuracy.

Sensitivity analysis across absolute thresholds from 27°C to 32°C confirmed that the 30°C threshold adopted a priori from Sakai et al. (2019) was near-optimal for both the >0% and ≥50% response definitions (Supplementary Table S3). AUC values formed a plateau between 29.5°C and 30.5°C, declining sharply at lower and higher thresholds. The ranking of metrics (threshold exceedance > DHW at ≥50%) was preserved under the alternative 8-week DHW accumulation window.

Abstract

4. Discussion

This study reports three descriptive findings from five years (2020–2024) of standardised coral bleaching monitoring across 26 sites and 585 survey points in the Ryukyu Archipelago and adjacent waters. First, bleaching prevalence is bimodally distributed, with a stable intermediate zone and inter-annual variation driven by redistribution between low and high domains. Second, the two major bleaching years in the dataset — 2022 and 2024 — are qualitatively distinct events that cannot be collapsed into a single "mass bleaching" category. Third, the relative performance of thermal stress metrics depends on the ecological severity of the response being predicted: threshold exceedance logic discriminates moderate-to-severe bleaching more effectively than cumulative DHW, but not bleaching at the detection limit — revealing that metric structure and response definition are inseparable considerations. These findings carry implications for how bleaching data are modelled, how bleaching events are classified, and how the design space of thermal stress metrics is explored.

4.1 Bimodality and threshold behaviour

The most striking structural feature of these data is the bimodal distribution of bleaching prevalence, confirmed independently by Hartigan's dip test (all years $p < 0.001$; Table 4) and beta mixture modelling (ΔBIC = 9.0–113.9; Table 5). The intermediate range (20–80%) accounts for a remarkably stable 21–27% of observations across all five years (Table 3), while inter-annual variation is driven almost entirely by redistribution between the low (<20%) and high (≥80%) domains. If bleaching progressed as a continuous gradient in response to thermal stress, one would expect the intermediate zone to expand in severe years; instead, it remains fixed. This pattern is consistent with a threshold-mediated binary switch rather than a dose-response continuum.

Multiple independent lines of physiological evidence support threshold behaviour in the bleaching response. Pfab et al. (2024) modelled the bioenergetic dynamics of symbiont photodamage and showed that once photodamage exceeds a critical level, energy supply to the CO_2-concentrating mechanism becomes insufficient, reactive oxygen species production accelerates, and symbiont expulsion becomes self-reinforcing — an escalating positive feedback that predicts minimal bleaching below the threshold and rapid progression to near-total bleaching above it. This model is situated within the broader mechanistic framework synthesised by Helgoe et al. (2024), who integrated multiple bleaching pathways into a triggers–cascades–endpoints architecture: diverse initial triggers (thermal, light, oxidative) converge on shared intracellular cascades that, once initiated, proceed largely independently of the triggering stimulus — a convergence structure that itself predicts threshold-like dynamics at the organismal level. Complementary to this ROS-centred cascade, Wooldridge (2009) proposed that thermal stress disrupts the CO_2-concentrating mechanism of Symbiodiniaceae, leading to photorespiratory production of glycolate and a toxic accumulation of hydrogen peroxide — a pathway that likewise predicts a sharp transition once the concentrating mechanism fails. These models differ in the specific rate-limiting step but converge on the prediction that bleaching initiation is governed by a threshold rather than by a graded dose-response. Our distributional data are consistent with this convergent prediction operating at the site level.

However, the distributional pattern alone cannot adjudicate between physiological threshold mechanisms and alternative explanations for bimodality. Spatial autocorrelation of environmental conditions — whereby neighbouring sites experience similar thermal, hydrodynamic, and light environments — could produce clustered bleaching outcomes that mimic threshold dynamics without requiring a physiological switch. Unmeasured covariates that vary at the site level (e.g., water flow regime, depth profile, species composition, symbiont community structure) could similarly generate bimodal outcomes through environmental filtering rather than through a discrete physiological transition. That the bimodal pattern is driven by between-site divergence rather than within-site variability (Section 3.5; ICC = 0.73–0.89) is necessary for but not sufficient to establish a physiological threshold mechanism: the same spatial polarisation would be expected under purely environmental sorting. More broadly, the same pattern can be expressed as a two-state mixture process at the network level, regardless of whether the underlying mechanism is physiological switching, environmental filtering, or their interaction. The bimodality is a robust descriptive finding; its causal attribution remains open.

That entire sites fell cleanly into either the low or the high domain in 2024 (Section 3.5) indicates that whatever determines bleaching outcome operates at the site scale or above. This spatial polarisation implies that site-level environmental factors — not captured by satellite SST alone — filter the bleaching response, whether through modulating a physiological threshold or through independent environmental determination.

To our knowledge, this bimodal pattern has not been explicitly reported in the coral bleaching literature as a persistent distributional feature across repeated annual surveys. Most large-scale analyses operate on bleaching occurrence as a binary response (Sully et al. 2019) or colony-level scores and cannot resolve the continuous site-level prevalence structure documented here. We note that Hughes et al. (2018b) described spatial heterogeneity in bleaching severity across the GBR, but their analysis focused on cumulative

impacts rather than on the distributional form of prevalence itself. The bimodal structure implies that zero-one inflated beta regression (Ospina & Ferrari 2012) is the appropriate modelling framework for these data, a point we return to in the Limitations.

If bimodality reflects a genuine state transition, it carries a structural implication for thermal stress prediction. The cumulative logic of DHW — in which bleaching intensity scales with integrated heat exposure — is designed to predict a graded dose-response. But if the ecologically relevant outcome is a binary switch between unbleached and bleached states, then what determines reef fate is whether a threshold is crossed, not how much heat is accumulated beyond it. This is not a problem that can be resolved by optimising DHW's internal parameters (accumulation window, HotSpot cutoff, baseline period); it is a question about whether the predictor category itself matches the outcome structure. The metric comparison in Section 3.7, where threshold exceedance logic outperforms DHW specifically for the state transition (≥50% bleaching) but not for detection-limit bleaching, is consistent with this interpretation.

4.2 Two anomalous years: 2022 versus 2024

Both 2022 and 2024 were major bleaching years, but their distributional signatures diverge in ways that carry interpretive consequences. In 2022, bleaching was positively skewed (skewness = 0.413): most sites experienced low-to-moderate bleaching, with a tail of severe cases. In 2024, the distribution was approximately symmetric (skewness = −0.100), the median reached 60.0%, and the high-bleaching component of the beta mixture model dominated (w■ = 0.616 vs 0.497). We characterise 2022 as a "partial mass bleaching" — selective, with a subset of sites severely affected — and 2024 as a "comprehensive mass bleaching" that engulfed the monitoring network.

This distinction matters because "mass bleaching" is routinely treated as a monolithic category. Existing classification schemes operate at different scales and capture different attributes: Spady et al. (2026) define global mass bleaching events by geographic extent across ocean basins, while Hughes et al. (2017) categorised aerial survey severity using an ordinal scale from no bleaching to extreme bleaching at the individual reef level. Neither framework captures the distributional structure of prevalence across a monitoring network — the within-event heterogeneity that distinguishes selective from comprehensive bleaching. We do not propose a typology from two events, but demonstrate that distributional parameters — specifically, mixture component weights and skewness — provide a descriptive vocabulary for this within-event variation that scalar summaries (mean prevalence, percent affected) and ordinal categories alike obscure. Whether partial and comprehensive bleaching represent discrete categories or endpoints of a continuum, and whether the distinction generalises beyond this dataset, are questions that require comparison across regions and monitoring programs. The descriptive vocabulary demonstrated here — partial versus comprehensive, skewed versus symmetric — is directly testable against other monitoring networks; if the same distributional signatures recur across independent datasets, a quantitative classification of mass bleaching events by their network-level distributional structure would become feasible.

Despite its greater extent, the 2024 event did not produce proportionally greater mortality at the time of survey (mean 18.7% vs 15.2% in 2022). Mellin et al. (2024) documented a similar decoupling between bleaching severity and coral cover change at continental scales on the GBR, suggesting that more extensive bleaching does not necessarily translate to greater structural loss.

As of our literature search (Web of Science, Scopus, Google Scholar; 6 February 2026), no peer-reviewed primary study has documented the 2024 bleaching event in Japan. Afzal et al. (2024) reported the 2022 event at Sekisei Lagoon. The present study provides the first systematic comparison of both events under identical monitoring protocols.

4.3 Bleaching, mortality, and the limits of annual monitoring

The relationship between bleaching prevalence and mortality is non-linear (Table 7). Mortality escalates sharply above 80% bleaching — mean 30.9%, with 96.4% of sites recording some mortality — while the 50–80% band shows mean mortality of only 13.2%, a 2.3-fold step across a single prevalence band. This threshold-like escalation in mortality is consistent with Brodnicke et al. (2019), who documented that severely bleached tabular *Acropora* colonies experienced markedly increased disease susceptibility and whole-colony mortality — a finding that supports the interpretation that high-severity bleaching fundamentally alters colony fate beyond the direct effects of thermal stress alone.

The most notable finding, however, is that even among sites with ≥80% bleaching, the median mortality was only 25.0%. Three-quarters of severely bleached sites had not experienced major mortality at the time of survey. As noted in Section 2.2, surveys are conducted 1–3 months after the bleaching peak, and Sakai et al. (2019) documented post-bleaching mortality progressing over six months with four-fold variation within the same genus (*Acropora*). The mortality values in Table 7 therefore represent a lower bound, and the true extent of post-bleaching mortality in 2022 and 2024 remains unknown.

Spot-level tracking provides further temporal context. The sharp decline in bleaching prevalence from 2022 (mean 37.6%) to 2023 (mean 12.3%) suggests rapid recovery, but annual snapshots cannot distinguish genuine tissue recovery from replacement by new recruits or from delayed mortality followed by recolonisation. Morais et al. (2021) distinguished two post-disturbance trajectories — boom-bust dynamics, in which fast-growing taxa (e.g., *Acropora*) dominate recovery but are disproportionately lost in subsequent disturbances, and protracted decline, in which massive taxa undergo slow, progressive erosion — a framework that underscores the interpretive limits of bleaching prevalence alone: an identical recovery trajectory in aggregate prevalence may mask divergent structural and demographic outcomes depending on which taxa drive the recovery. Back-to-back tracking across 2022 and 2024 reveals that 32.0% of survey points experienced ≥50% bleaching in both years — chronic high-stress sites that either failed to recover fully or were re-bleached — while 9.4% transitioned from unaffected in 2022 to severely bleached in 2024. Whether this reflects cumulative damage or the loss of a site-level protective factor cannot be resolved from bleaching prevalence data alone; cross-referencing with coral cover change is required.

4.4 Metric structure, response definition, and the design space of thermal stress metrics

The predictive comparison between threshold exceedance days and DHW reveals that relative performance depends on the response definition. For bleaching at the detection limit (>0%), DHW (AUC 0.761) and days above 30°C (AUC 0.748) performed comparably (DeLong test, p = 0.214). For moderate bleaching (≥50%), threshold exceedance was significantly superior (AUC 0.846 vs 0.758; p < 0.001). For severe bleaching (≥80%), both metrics performed well with no statistically significant difference (AUC 0.854 vs 0.843; p = 0.299).

This response-definition dependence has a coherent interpretation. The >0% detection-limit threshold conflates near-noise signals — including observer variability and transient paling at the recording resolution of ~0.5% — with ecologically meaningful bleaching. DHW, as a cumulative metric, may capture broad heat exposure conducive to any detectable paling, whereas threshold exceedance logic better captures the transition to sustained, ecologically consequential bleaching. The threshold metric's AUC jump from 0.748 (>0%) to 0.846 (≥50%) directly mirrors the binary dynamics of the bimodal distribution itself: the ecologically significant state transition documented in Section 3.2 is, by definition, a transition that occurs above, not at, the detection limit.

This reframes the question from "which metric is better?" to "which metric is appropriate for which ecological question?" DHW retains utility as a predictor of broad thermal stress exposure — the onset of any detectable bleaching signal. Threshold exceedance logic is better suited to predicting whether a site will undergo ecologically significant bleaching, the state transition that determines reef trajectory.

This response-definition dependence has a structural parallel in psychiatric epidemiology, where prevalence estimates and risk factor effect sizes vary systematically with the diagnostic threshold applied — a problem that has driven two decades of debate over dimensional versus categorical classification in psychopathology, and that has never been resolved by optimising a single threshold. The parallel is not merely analogical: in both domains, the outcome measure is a continuous variable dichotomised at an arbitrary threshold, and the choice of that threshold determines which predictor appears superior. The implication is methodological: rather than reporting AUC at a single response definition, metric comparison studies should report AUC as a function of response threshold — a profile that makes the dependence visible rather than assuming it away.

However, a fundamental limitation of the days-above-30°C metric must be acknowledged. The monitoring network spans 6° of latitude, and the 30°C threshold has an inherent geographic bias: at high-latitude sites (e.g., Kushimoto, 33.5°N), summer SST rarely reaches 30°C even in anomalously warm years, yet bleaching does occur. In these locations, the metric yields zero days above 30°C regardless of thermal stress intensity, because the absolute threshold is set above the site's climatological range. DHW, by contrast, uses site-specific MMM-relative anomalies that automatically adjust for latitudinal differences in baseline temperature — a structural advantage that explains why DHW matches or exceeds threshold exceedance performance at the >0% response level.

This analysis points toward a natural direction for metric development. The predictive advantage of threshold exceedance logic for ecologically significant bleaching suggests that *whether a critical temperature is exceeded* matters more than *how much cumulative heat is absorbed*. But the failure of a single absolute threshold across the latitudinal range demonstrates that the critical temperature itself must be site-specific. The logical endpoint is a site-specific MMM-relative threshold exceedance count — a metric that combines DHW's climatological relativism with threshold exceedance logic's binary structure. Such a metric would ask: on how many days did SST exceed MMM + Δ°C, where Δ is optimised regionally or locally? This formulation subsumes both DHW ($\Delta = 0$, cumulative) and absolute threshold counting (Δ = fixed offset from a universal baseline) as special cases, and represents a logical next step in exploring the design space of thermal stress metrics. Cheung et al. (2025) demonstrated through interpretable machine learning on the GBR that non-linear interactions between heat stress and environmental covariates — including water flow, light availability, and thermal history — resolve bleaching outcomes that accumulated heat stress alone cannot predict, reinforcing the case for moving beyond single-variable cumulative metrics toward context-dependent predictive frameworks.

Prior DHW optimisation studies have focused on improving internal parameters — accumulation window length (Lachs et al. 2021), cutoff temperature and regional calibration (Whitaker & DeCarlo 2024). Sensitivity analysis across thresholds from 27°C to 32°C confirmed that 30°C is near-optimal for this network (Supplementary Table S3), with AUC forming a plateau between 29.5°C and 30.5°C — a zone rather than a point. Our findings extend this parameter-tuning approach by raising a structural question about the design space: whether cumulative integration is the appropriate predictor category for ecologically significant bleaching, and whether the transition from cumulative to threshold logic entails a simultaneous transition from absolute to relative thresholds. This is not a challenge to DHW's operational framework, but an exploration of the broader design space within which DHW occupies a specific — and for many applications, appropriate — position.

4.5 Limitations

Six limitations qualify these findings.

First, surveys are annual autumn snapshots (Section 2.2). The mortality values reported here are conservative estimates that do not capture the full delayed mortality process.

Second, vertical and lateral environmental data — water flow, turbidity, light environment, dissolved oxygen, temperature stratification — are absent from the monitoring protocol. The site-level environmental factors that produce the spatial polarisation documented in Section 4.1 remain unidentified. Without direct environmental measurements, the environmental heterogeneity hypothesis is observationally motivated but untested.

Third, the 2021 Yaeyama region presents a pattern consistent with non-random logger deployment: logger sites recorded near-zero bleaching while non-logger sites recorded 100% bleaching (Fisher's exact test, p ■ 0.001; detailed cross-tabulation in the Supplementary Material, Table S2). The monitoring protocol directs loggers to locations representative of the survey site and where temperature fluctuations are pronounced, criteria that may favour microhabitats where coral communities have historically persisted — a form of non-random selection into observation analogous to Missing Not At Random conditions (Little & Rubin 2019; Bowler et al. 2024). If logger sites systematically sample cooler or more sheltered microhabitats, satellite-logger temperature discrepancies may overestimate satellite bias. This single-region, single-year observation does not establish the generality or magnitude of the bias, but it signals a structural limitation in logger-based ground-truthing that warrants systematic investigation.

Fourth, species-level resolution is lacking. Bleaching prevalence is a species-agnostic aggregate, and the mortality decoupling documented in Section 4.3 is likely driven by species, genotype, and symbiont differences that site-level surveys cannot resolve. Species composition provides an alternative explanation for bimodality: sites dominated by thermally sensitive genera (e.g., *Acropora*) may cluster in the high-bleaching domain independent of any physiological threshold mechanism. The Monitoring Site 1000 protocol records genus-level composition at a subset of sites; systematic analysis of composition effects on the distributional structure is deferred to subsequent work.

Fifth, spatial autocorrelation is not accounted for. Survey points within sites are spatially proximate and not independent. The descriptive statistics reported here do not adjust for this non-independence; mixed-effects models with site as a random effect are required for inferential analyses. Sample sizes also varied across years (428–496), driven by site additions and removals (Table 1); sensitivity analyses on the balanced panel of 21 sites are reported in the Supplementary Material.

Sixth, the predictive comparison (Section 3.7) reveals that metric performance is sensitive to the binary response threshold. The AUC values in Table 8 show that the advantage of threshold exceedance over DHW is statistically significant only for moderate bleaching ($\geq$50%; DeLong test, $p < 0.001$), with no significant difference at the detection limit ($>$0%; $p = 0.214$) or at the severe bleaching threshold ($\geq$80%; $p = 0.299$). This response-definition dependence means that claims about metric superiority cannot be generalised without specifying the ecological severity of the outcome being predicted. Furthermore, any binary classification of a continuous variable entails information loss; the zero-one inflated beta regression framework identified in Section 4.1 would allow modelling the full distributional response without dichotomisation. The DeLong test assumes independence of observations; the repeated-measures structure of this dataset (multiple years per site) means that the reported p-values should be interpreted with caution, and mixed-effects logistic regression or GEE would be required for fully corrected inferential tests.

The bimodal distributional structure implies that standard linear or beta regression models are misspecified for these data; zero-one inflated beta regression (Ospina & Ferrari 2012) is the statistically appropriate framework, and its implementation is a priority for subsequent work.

What this study establishes and what remains open

Established. (1) Bleaching prevalence in the Ryukyu Archipelago is bimodally distributed, with a stable intermediate zone — a pattern consistent with threshold-mediated state transitions rather than continuous dose-response. (2) The 2022 and 2024 bleaching events are qualitatively distinct: partial versus comprehensive mass bleaching, distinguishable by distributional parameters that provide a descriptive

vocabulary beyond scalar summaries. (3) The predictive performance of thermal stress metrics depends on response definition: threshold exceedance logic discriminates moderate-to-severe bleaching ($\geq 50\%$) more effectively than cumulative DHW (DeLong test, $p < 0.001$), while both perform comparably for bleaching at the detection limit ($p = 0.214$) — indicating that metric structure and ecological severity threshold are inseparable design considerations.

Open. (1) The environmental factors that produce site-level spatial polarisation. (2) The generality and magnitude of logger deployment bias across the monitoring network. (3) Whether the bimodality reflects physiological thresholds, environmental filtering, or their interaction. (4) The full extent of post-bleaching mortality beyond the autumn survey window. (5) Optimal site-specific MMM-relative threshold values across the latitudinal range of the monitoring network. (6) Whether a site-specific MMM-relative threshold exceedance count — combining DHW's climatological relativism with threshold exceedance logic — improves predictive accuracy across response definitions and latitudinal gradients.

Figure caption

Fig. 1. Annual distribution of coral bleaching prevalence (upper panel) and mortality prevalence (lower panel) across the Monitoring Site 1000 network, fiscal years 2020–2024. Violin plots show kernel density estimates (Gaussian kernel, Scott's rule bandwidth); KDE tails may extend slightly beyond the 0–100% data bounds. Box plots within each violin show median (horizontal line) and interquartile range. The percentage of survey points with bleaching > 0% (upper panel) or mortality > 0% (lower panel) is annotated above each year. Sample sizes: n = 428, 438, 496, 455, 470 for FY2020–2024, respectively.

Abstract

References

Afzal MS, Uehara H, Yamashiro H (2024) Recurrent mass bleaching of coral assemblages in the Sekisei Lagoon, Japan. *Galaxea, Journal of Coral Reef Studies* 26: 1–9.

Bowler DE, Boyd RJ, Callaghan CT, Robinson RA, Isaac NJB, Pocock MJO (2024) Treating gaps and biases in biodiversity data as a missing data problem. *Biological Reviews* 100(1): 50–67. doi:10.1111/brv.13127

Brodnicke OB, Bourne DG, Heron SF, Pears RJ, Stella JS, Smith HA, Willis BL (2019) Unravelling the links between heat stress, bleaching and disease: fate of tabular corals following a combined disease and bleaching event. *Coral Reefs* 38: 591–603. doi:10.1007/s00338-019-01813-9

Cheung MWM, Chaloupka M, Hock K, Mumby PJ (2025) Moving beyond temperature metrics in coral bleaching prediction using interpretable machine learning. *Global Ecology and Biogeography* 34: e70105. doi:10.1111/geb.70105

DeLong ER, DeLong DM, Clarke-Pearson DL (1988) Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics* 44(3): 837–845.

Hartigan JA, Hartigan PM (1985) The dip test of unimodality. *Annals of Statistics* 13(1): 70–84.

Helgoe J, Davy SK, Weis VM, Kemp DW (2024) Triggers, cascades, and endpoints: connecting the dots of coral bleaching mechanisms. *Biological Reviews* 99(3): 715–752. doi:10.1111/brv.13042

Hughes TP, Kerry JT, Álvarez-Noriega M, Álvarez-Romero JG, Anderson KD, Baird AH, Babcock RC, Beger M, Bellwood DR, Berkelmans R, Bridge TC, Butler IR, Byrne M, Cantin NE, Comeau S, Connolly SR, Cumming GS, Dalton SJ, Diaz-Pulido G, Eakin CM, Figueira WF, Gilmour JP, Harrison HB, Heron SF, Hoey AS, Hobbs J-PA, Hoogenboom MO, Kennedy EV, Kuo C-Y, Lough JM, Lowe RJ, Liu G,

McCulloch MT, Malcolm HA, McWilliam MJ, Pandolfi JM, Pears RJ, Pratchett MS, Schoepf V, Simpson T, Skirving WJ, Sommer B, Torda G, Wachenfeld DR, Willis BL, Wilson SK (2017) Global warming and recurrent mass bleaching of corals. *Nature* 543: 373–377. doi:10.1038/nature21707

Hughes TP, Anderson KD, Connolly SR, Heron SF, Kerry JT, Lough JM, Baird AH, Baum JK, Berumen ML, Bridge TC, Claar DC, Eakin CM, Gilmour JP, Graham NAJ, Harrison HB, Hobbs J-PA, Hoey AS, Hoogenboom M, Lowe RJ, McCulloch MT, Pandolfi JM, Pratchett MS, Schoepf V, Torda G, Wilson SK (2018a) Spatial and temporal patterns of mass bleaching of corals in the Anthropocene. *Science* 359: 80–83. doi:10.1126/science.aan8048

Hughes TP, Kerry JT, Baird AH, Connolly SR, Dietzel A, Eakin CM, Heron SF, Hoey AS, Hoogenboom MO, Liu G, McWilliam MJ, Pears RJ, Pratchett MS, Skirving WJ, Stella JS, Torda G (2018b) Global warming transforms coral reef assemblages. *Nature* 556: 492–496. doi:10.1038/s41586-018-0041-2

JPL MUR MEaSUREs Project (2015) GHRSST Level 4 MUR Global Foundation Sea Surface Temperature Analysis (v4.1). doi:10.5067/GHGMR-4FJ04

Kass RE, Raftery AE (1995) Bayes factors. *Journal of the American Statistical Association* 90(430): 773–795.

Kawagoe S (2017) Mass coral bleaching in 2016 reported by the Monitoring sites 1000 project. *Journal of Japanese Coral Reef Society* 19: 21–28. [in Japanese]

Kimura T, Chou LM, Huang D, Tun K, Goh E (eds) (2022) Status and Trends of East Asian Coral Reefs: 1983–2019. Global Coral Reef Monitoring Network (GCRMN) / International Coral Reef Initiative (ICRI).

Lachs L, Bythell JC, East HK, Edwards AJ, Mumby PJ, Skirving WJ, Spady BL, Guest JR (2021) Fine-tuning heat stress algorithms to optimise global predictions of mass coral bleaching. *Remote Sensing* 13(14): 2677. doi:10.3390/rs13142677

Little RJA, Rubin DB (2019) *Statistical Analysis with Missing Data*, 3rd ed. John Wiley & Sons.

Liu G, Heron SF, Eakin CM, Muller-Karger FE, Vega-Rodriguez M, Guild LS, De La Cour JL, Geiger EF, Skirving WJ, Burgess TFR, Strong AE, Harris A, Maturi E, Ignatov A, Sapper J, Li J, Lynds S (2014) Reef-scale thermal stress monitoring of coral ecosystems: new 5-km global products from NOAA Coral Reef Watch. *Remote Sensing* 6(11): 11579–11606. doi:10.3390/rs61111579

Mellin C, Matthews S, Anthony KRN, Brown SC, Caley MJ, Johns KA, Osborne K, Puotinen M, Thompson A, Wolff NH, Fordham DA, MacNeil MA (2024) Spatial resilience of the Great Barrier Reef under cumulative disturbance impacts. *Proceedings of the Royal Society B* 291: 20240095.

Ministry of Environment, Japan (2024) モニタリングサイト1000　サンゴ礁調査マニュアル [Monitoring Site 1000 Coral Reef Survey Manual]. https://www.biodic.go.jp/moni1000/manual/coral_manual.pdf

Morais RALS, Connolly SR, Bellwood DR (2021) Dangerous demographics in post-bleach corals reveal boom-bust versus protracted declines. *Scientific Reports* 11: 18787. doi:10.1038/s41598-021-98239-7

Ospina R, Ferrari SLP (2012) A general class of zero-or-one inflated beta regression models. *Computational Statistics & Data Analysis* 56(6): 1609–1623. doi:10.1016/j.csda.2011.10.005

Pfab F, Detmer AR, Engel V, Kayal M, Donovan MK, Edmunds PJ, Nisbet RM (2024) A bioenergetic model for coral bleaching. *Coral Reefs* 43: 1345–1358.

Sakai K, Singh T, Iguchi A (2019) Bleaching and post-bleaching mortality of *Acropora* corals on a heat-susceptible reef in 2016. *PeerJ* 7: e8138. doi:10.7717/peerj.8138

Scott DW (1992) *Multivariate Density Estimation: Theory, Practice, and Visualization*. John Wiley & Sons.

Spady BL, Skirving WJ, De La Cour JL, Geiger EF, Liu G, Hoegh-Guldberg O, Norrie A, Heron SF, Pomeroy MW, Kolodziej G, Brown KT, Manzello D (2026) The 4th global coral bleaching event: ushering in an era of near-annual bleaching. *Coral Reefs*. doi:10.1007/s00338-025-02810-x

Sully S, Burkepile DE, Donovan MK, Hodgson G, van Woesik R (2019) A global analysis of coral bleaching over the past two decades. *Nature Communications* 10: 1264. doi:10.1038/s41467-019-09238-2

Whitaker H, DeCarlo TM (2024) Re(de)fining degree-heating week: coral bleaching variability necessitates regional and temporal optimization of global forecast model stress metrics. *Coral Reefs* 43: 969–984. doi:10.1007/s00338-024-02512-w

Wooldridge SA (2009) A new conceptual model for the warm-water breakdown of the coral–algae endosymbiosis. *Marine and Freshwater Research* 60: 483–496. doi:10.1071/MF08251