

Northumbria Research Link

Citation: Park, Sung-Hong, Leka, K. D., Kusano, Kanya, Andries, Jesse, Barnes, Graham, Bingham, Suzy, Bloomfield, Shaun, McCloskey, Aoife E., Delouille, Veronique, Falconer, David, Gallagher, Peter T., Georgoulis, Manolis K., Kubo, Yuki, Lee, Kangjin, Lee, Sangwoo, Lobzin, Vasily, Mun, JunChul, Murray, Sophie A., Hamad Nageem, Tarek A. M., Qahwaji, Rami, Sharpe, Michael, Steenburgh, R. A., Steward, Graham and Terkildsen, Michael (2020) A Comparison of Flare Forecasting Methods. IV. Evaluating Consecutive-day Forecasting Patterns. *The Astrophysical Journal*, 890 (2). p. 124. ISSN 1538-4357

Published by: The American Astronomical Society

URL: <https://doi.org/10.3847/1538-4357/ab65f0> <<https://doi.org/10.3847/1538-4357/ab65f0>>

This version was downloaded from Northumbria Research Link:
<http://nrl.northumbria.ac.uk/id/eprint/42453/>

Northumbria University has developed Northumbria Research Link (NRL) to enable users to access the University's research output. Copyright © and moral rights for items on NRL are retained by the individual author(s) and/or other copyright owners. Single copies of full items can be reproduced, displayed or performed, and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided the authors, title and full bibliographic details are given, as well as a hyperlink and/or URL to the original metadata page. The content must not be changed in any way. Full items must not be sold commercially in any format or medium without formal permission of the copyright holder. The full policy is available online: <http://nrl.northumbria.ac.uk/policies.html>

This document may differ from the final, published version of the research and has been made available online in accordance with publisher policies. To read and/or cite from the published version of the research, please visit the publisher's website (a subscription may be required.)



A Comparison of Flare Forecasting Methods. IV. Evaluating Consecutive-day Forecasting Patterns

Sung-Hong Park¹, K. D. Leka^{1,2}, Kanya Kusano¹, Jesse Andries³, Graham Barnes², Suzy Bingham⁴, D. Shaun Bloomfield⁵, Aoife E. McCloskey⁶, Veronique Delouille³, David Falconer⁷, Peter T. Gallagher⁸, Manolis K. Georgoulis^{9,10}, Yuki Kubo¹¹, Kangjin Lee^{12,13}, Sangwoo Lee¹⁴, Vasily Lobzin¹⁵, JunChul Mun¹⁶, Sophie A. Murray^{6,8}, Tarek A. M. Hamad Nageem¹⁷, Rami Qahwaji¹⁷, Michael Sharpe⁴, R. A. Steenburgh¹⁸, Graham Stewart¹⁵, and Michael Terkildsen¹⁵

¹ Institute for Space-Earth Environmental Research, Nagoya University, Nagoya, Japan; shpark@isee.nagoya-u.ac.jp

² NorthWest Research Associates, Boulder, CO, USA

³ Solar-Terrestrial Center for Excellence, Royal Observatory of Belgium, Brussels, Belgium

⁴ Met Office, Exeter, UK

⁵ Northumbria University, Newcastle upon Tyne, UK

⁶ School of Physics, Trinity College Dublin, Dublin, Ireland

⁷ NASA/NSSTC, Huntsville, AL, USA

⁸ School of Cosmic Physics, Dublin Institute for Advanced Studies, Dublin, Ireland

⁹ Department of Physics & Astronomy, Georgia State University, Atlanta, GA, USA

¹⁰ Research Center Astronomy and Applied Mathematics, Academy of Athens, Athens, Greece

¹¹ National Institute of Information and Communications Technology, Tokyo, Japan

¹² Electronics and Telecommunications Research Institute, Daejeon, Republic of Korea

¹³ School of Space Research, Kyung Hee University, Yongin, Republic of Korea

¹⁴ SELab, Inc., Seoul, Republic of Korea

¹⁵ Space Weather Services, Bureau of Meteorology, Sydney, Australia

¹⁶ Korean Space Weather Center, National Radio Research Agency, Jeju, Republic of Korea

¹⁷ University of Bradford, Bradford, UK

¹⁸ NOAA/NWS/NCEP Space Weather Prediction Center, Boulder, CO, USA

Received 2019 November 11; revised 2019 December 18; accepted 2019 December 26; published 2020 February 19

Abstract

A crucial challenge to successful flare prediction is forecasting periods that transition between “flare-quiet” and “flare-active.” Building on earlier studies in this series in which we describe the methodology, details, and results of flare forecasting comparison efforts, we focus here on patterns of forecast outcomes (success and failure) over multiday periods. A novel analysis is developed to evaluate forecasting success in the context of catching the first event of flare-active periods and, conversely, correctly predicting declining flare activity. We demonstrate these evaluation methods graphically and quantitatively as they provide both quick comparative evaluations and options for detailed analysis. For the testing interval 2016–2017, we determine the relative frequency distribution of two-day dichotomous forecast outcomes for three different event histories (i.e., event/event, no-event/event, and event/no-event) and use it to highlight performance differences between forecasting methods. A trend is identified across all forecasting methods that a high/low forecast probability on day 1 remains high/low on day 2, even though flaring activity is transitioning. For M-class and larger flares, we find that explicitly including persistence or prior flare history in computing forecasts helps to improve overall forecast performance. It is also found that using magnetic/modern data leads to improvement in catching the first-event/first-no-event transitions. Finally, 15% of major (i.e., M-class or above) flare days over the testing interval were effectively missed due to a lack of observations from instruments away from the Earth–Sun line.

Unified Astronomy Thesaurus concepts: [The Sun \(1693\)](#); [Solar magnetic fields \(1503\)](#); [Solar flares \(1496\)](#); [Solar activity \(1475\)](#); [Solar active region magnetic fields \(1975\)](#); [Sunspots \(1653\)](#); [Solar x-ray flares \(1816\)](#); [Astronomy data analysis \(1858\)](#); [Astronomy data visualization \(1968\)](#); [Astrostatistics tools \(1887\)](#)

1. Introduction

Forecasting solar flares provides a laboratory with which to examine the understanding of these energetic events, but forecasts also serve to protect infrastructure impacted by our Sun’s variable output on a daily basis. The success of forecasting, from both an empirical and a physical point of view, has thus far been measured using statistical evaluations of correct forecasts with each event considered independently and equally. Both operationally and physically, however, it is crucial to understand the transitions from “flare-quiet” to “flare-active” and back again, as the Sun and its magnetic fields evolve, generating, storing, and finally releasing free magnetic energy in the form of energetic events.

A focused workshop on “Benchmarks for Operational Solar Flare Forecasting Systems” was held in 2017 at the Institute for Space-Earth Environmental Research (ISEE), Nagoya University. The primary objective of the workshop was to compare in a quantitative manner the performance characteristics of today’s operational flare forecasting methods as a follow-up to the “All Clear” workshop and its initial investigation into the methodology of forecast comparisons (Barnes et al. 2016, hereafter Paper I). For this workshop, forecasts from 19 different operational flare forecasting methods were submitted for an agreed-upon testing interval of 2016 January 1 to 2017 December 31, following agreed-upon forecast intervals and event definitions described in

Leka et al. (2019a, hereafter Paper II). Results focusing on the head-to-head comparisons are presented there using multiple evaluation methodologies, including graphics and quantitative metrics based on both probabilistic and dichotomous forecasts, as are often used for forecast validation (Woodcock 1976; Bloomfield et al. 2012; Barnes et al. 2016; Kubo et al. 2017; Murray et al. 2017, and references therein). Recognizing the small sample size and short testing period, it was found in Paper II that (1) many methods consistently demonstrate skill, although (2) no single method is “best” across multiple metrics, and (3) no method performs “well” (i.e., better than 0.5 across numerous normalized skill scores and validation metrics, where 0.5 is halfway between no skill and perfect). Most importantly, the required methodology for providing fair and meaningful comparisons across forecasts was demonstrated, centering primarily on common testing intervals, event definitions, and evaluation using a variety of metrics. The question of *why* certain methods performed better or worse than others was examined in Leka et al. (2019b, hereafter Paper III) by means of grouping the methods in different categories according to their implementation details. In this context of broad implementation differences, the behavior and performance of the methods were evaluated. The results were weak due to both the nonuniqueness of the categorizations and the small sample size, but it was found that including prior flare history and active region evolution likely led to improved performance, with a further indication that including a human “forecaster in the loop” (FITL) was also advantageous.

During the workshop, the participants expressed interest in examining a particular interval in detail, e.g., a case study, in part due to the fact that NOAA active region (AR) 12673 was fresh in our memories, having produced at least one flare greater than or equal to *GOES* M1.0 level each day for 7 consecutive days from 2017 September 4. In a cursory manner, we found that many of the methods failed to predict a high chance of major flares for the first day of AR 12673’s multiday flaring activity. Yet while some methods subsequently and significantly increased their forecast probabilities on the second day so that they successfully predicted the second day of activity, other methods’ forecasts showed little change for that second day (i.e., the event day was missed again) regardless of the large flares that occurred on the previous day. These different behaviors between forecasting methods motivated us to explore forecast performance over consecutive days with variable event histories.

Case studies are often used by operational facilities during forecaster training to target a particular known failure. Questions often asked in case-study examinations include “was the first flare (of a series over a multiday period) predicted correctly?” and “did forecast probabilities in fact decrease as flare activity subsided?” Such case studies can be misleading, however, as a method’s performance during a particular interval may not reflect its performance when numerous different intervals are considered. Here we extend this line of questioning to examine particular patterns of forecasting behavior using a multiday analysis; specifically, we examine sets of 2 consecutive days where at least one of those days includes an event. We test the hypothesis that including some aspect of prior behavior or temporal evolution results in forecasts that are able to better adjust for varying flare activity. To evaluate this hypothesis, we present a newly developed analysis methodology to quantitatively evaluate specific temporally oriented performance characteristics of solar flare forecasts.

2. Methodology

We describe here the input data from the participating methods and the methodology employed for evaluating outcome patterns of daily forecasts over consecutive forecast days, i.e., in the context of the challenges of predicting the first flaring and flare-quiet days described above. The results are presented later in Section 3.

2.1. Participating Flare Forecasting Methods

The participants of the ISEE workshop brought 19 operational flare forecasting methods for analysis. Among them are several methods that have been implemented as operational flare forecasting systems at space weather Regional Warning Centres (RWCs), as well as at research institutions. Note that while no human forecaster intervenes in the forecast output of any research institution-based methods, in general there are experienced forecasters at RWCs who take into account the implemented method outputs and may adjust them prior to issuing their official forecasts. Details of all participating flare forecasting methods can be found in Paper II and references therein. As in Paper III, for reference, we reproduce an abbreviated version in Appendix A as Table 7, which lists the methods, relevant publications, and monikers/acronyms used here. As is clear from the earlier papers and Table 7 here, not all of the 19 methods are completely independent; in some cases, they consist of different implementations of the same general approach (e.g., the four versions of MAG4).

Full-disk daily forecasts were submitted and processed such that two event definitions were used consistently: 24 hr validity periods, effectively zero-hour latencies, but then with two different lower limits of C1.0 and M1.0 in the *GOES* flare class; these are referred to here as C1.0+/0/24 and M1.0+/0/24, respectively. For the methods that did not produce such full-disk exceedance forecasts, we converted their region-based forecasts to full-disk forecasts (described in Appendix B.1. of Paper II) and, when appropriate, combined category-limited (i.e., C1.0–C9.9, M1.0–M9.9, and X1.0+) forecasts using conditional probabilities to provide exceedance forecasts (discussed in Appendix B.2 of Paper II). The testing interval was 2016 January 1 to 2017 December 31, inclusive. For each method and event definition, a binary (yes/no) event list is compiled from recorded *GOES* flares in the NOAA Edited Solar Event Lists.¹⁹ Most of the methods issue their forecasts at 00:00 UT, while SIDC issues at 12:30 UT and NICT issues at 06:00 UT; for the latter two methods, custom event lists were created for the most consistent comparison possible with the other methods.

Over the course of the 731 days in 2016–2017, there are 188 event days (25.7%) and 26 event days (3.6%) for C1.0+/0/24 and M1.0+/0/24, respectively. In the case of missing forecasts for a method, these days are filled with 0% probability values; as discussed in Paper II, this is detrimental to the performance evaluation but is fair for operational purposes. Probabilistic forecasts can be converted into dichotomous forecasts by setting a threshold P_{th} above which a forecast is classified as forecasting positively for an event. In this study, two different values of P_{th} are used. As with Papers II and III, $P_{th} = 0.5$ is applied by default, but in addition, we examine the impact on performance when the threshold reflects the testing

¹⁹ <ftp://ftp.swpc.noaa.gov/pub/warehouse>




Figure 1. An example of forecast probabilities (solid line) from an anonymous flare forecasting method for the $M1.0+0/24$ event definition is shown with the highest *GOES* soft X-ray flux (dashed line) observed each day over the interval of 2017 September 1–15. The *GOES* M1.0 level is marked with a horizontal dotted line for reference. Each day’s forecast has a color-coded background shading that indicates one of the resulting dichotomous forecast outcomes using a $P_{th} = 0.5$ level: correct null (gray), miss (red), hit (green), and false alarm (blue).

interval climatology instead: $P_{th} = 0.257$ and 0.036 for $C1.0+0/24$ and $M1.0+0/24$, respectively. Note that this event frequency is very low (as discussed in Paper II), reflecting the fact that our testing period occurs on the declining phase of a weak-activity solar cycle. All submitted forecast probabilities, as well as relevant codes for data processing and analysis, are freely available (Leka & Park 2019) so that readers may explore the metrics and effects of varying P_{th} as desired.

2.2. At-a-glance Performance

We first investigate the overall performance of the daily flare forecasts from the 19 methods under consideration using a color-coded diagram of the forecasts in dichotomous form, as demonstrated in Figure 1. The diagram uses a designated P_{th} to color-code the dichotomous forecast outcomes (i.e., hits, correct nulls, false alarms, misses), and the daily highest *GOES* soft X-ray flux is shown as well. The results are discussed in Section 3.1.

2.3. Two-day Analysis

Examining consecutive-day forecasting patterns enables us to begin a statistical analysis of that which is of interest in case studies. We consider three different “event histories” when one or both of the days includes an event (see Table 1); for simplicity, we do not consider the no-event/no-event history. The provided forecasts for each of the 2 days then produce a “forecast outcome pattern” with four possible outcomes. That is, the two forecasts for the two-day period are considered as a unit (as opposed to each flare event being considered independently).

A goal here is to highlight misforecasting patterns in the context of the “first event” and “first quiet” (effectively the “last event”) of a flare-active period. In this context, the two event-history options are of an event occurring after a period of quiet (such as when a region begins to be flare-active) during times of high but possibly varied flare activity and no event occurring after a flare-active period, meaning in this context the first flare-quiet day when activity is diminishing.

For this analysis, we extend the graphical summary from the “at-a-glance” diagram in Figure 1 to the categorizations in Table 1 and use a radar-plot format to summarize the performance characteristics of consecutive-day forecasts (Figure 2). For a specified P_{th} , we compute the relative frequency with which a method’s forecasts fall under each possible outcome pattern. For example, the number of occurrences of a particular outcome pattern (e.g., the number of C-H outcomes) is divided by the total number of two-day forecasts within that particular event-day history (in this case, no-event/event) over the 2 yr interval (or 13 for $M1.0+0/24$ and 66 for $C1.0+0/24$). Examples of radar plots and how they display particular outcomes (perfect, systematically over-forecasting, etc.) are shown in Figure 2. This presentation method statistically summarizes some of the important points of case studies.

For the specific question of how well the methods predict the first flare/first quiet (which is more explicitly the correct prediction of a change in activity from quiet to flare-active and a change in activity from flare-active back to quiet), we can additionally specify which of the mixed-event outcome patterns have more or less impact on overall forecast performance. Successfully forecasting the first flaring day requires that at best, both the no-event and the following event day are correctly forecast; if only one of the two days is correctly forecast, it should at least be the event day rather than the no-event day. In other words, focusing on the no-event/event history, good forecasting performance dictates that the “two-day-correct” outcomes exceed the “two-day-incorrect” outcomes and the “first-day-incorrect/second-day-correct” outcomes exceed the “first-day-correct/second-day-incorrect” outcomes, if using the labels $C-H > F-M$ and $F-H > C-M$. For the radar plots, this translates to asymmetric relative frequencies across particular 180° sectors, i.e., on an analog clock face, $11:00 > 05:00$ and $02:00 > 08:00$, respectively. Conversely, better performance forecasting of the first flare-quiet day focuses on the event/no-event history and requires $H-C > M-F$ ($01:00 > 07:00$) and $M-C > H-F$ ($04:00 > 10:00$).

2.4. Two-day Analysis Plus Categorization

Finally, we ask what implementation factors contribute to performance for consecutive-day forecasts. In this context, we examine our original hypothesis that explicitly including temporal information in the forecasting method would improve performance, including an improved ability to catch the first-event/first-no-event transitions. To this end, we focus on a few of the broad implementation options adopted in Paper III (see Table 5 in Paper III for the assignment of each forecasting method according to implementation option) and group the results by the outcome patterns. In some cases, three options as presented in Paper III are reduced to binary options here, as indicated, in order to maximize the sample size. The binary implementation options (BIOs) that we focus on here are as follows.

The *training interval* describes a method’s training data as “Short,” “Long,” or “Hybrid,” which generally correspond to *Solar Dynamics Observatory* (*SDO*; Pesnell et al. 2012) data only, multi-solar-cycle training, or a combination (e.g., encompassing longer training but then using *SDO* data for the forecasts themselves), respectively. For this work, we group the *Short* and *Hybrid* options together.

Table 1
Outcome Pattern Summary for Consecutive-day Forecast Analysis

Event History Day 1/Day 2	If Forecast Is Day 1/Day 2	Then Outcome Is Day 1/Day 2	Label	No. of Instances (% of the Total)	
				C1.0+/0/24	M1.0+/0/24
Event/Event	Yes/Yes	Hit/Hit	H-H	121 (16.6%)	12 (1.6%)
	Yes/No	Hit/Miss	H-M		
	No/Yes	Miss/Hit	M-H		
	No/No	Miss/Miss	M-M		
No Event/Event	Yes/Yes	False alarm/Hit	F-H	66 (9.0%)	13 (1.8%)
	Yes/No	False alarm/Miss	F-M		
	No/Yes	Correct null/Hit	C-H		
	No/No	Correct null/Miss	C-M		
Event/No Event	Yes/Yes	Hit/False alarm	H-F	67 (9.2%)	14 (1.9%)
	Yes/No	Hit/Correct null	H-C		
	No/Yes	Miss/False alarm	M-F		
	No/No	Miss/Correct null	M-C		

Data characterization divides the methods into two broad groups: “*Simple*,” which relies on qualitative analysis or simpler inputs (e.g., sunspot group categories), or “*Magnetic/Modern*” quantitative analysis, most often comprised of photospheric magnetic field data.

Persistence or prior flare activity describes whether a method qualitatively or quantitatively included persistence or prior flare activity in computing forecasts (or, alternatively, that no such information was included).

The *evolution* of underlying active regions is included explicitly in some methods, implicitly in others, but not at all for the majority of methods. Including evolution could take the form of, for example, tracking the evolution of the sunspot group class and its impact on flaring rates (as for MCEVOL; see McCloskey et al. 2016) or the contribution of an FITL in judging a perceived change in a region’s flaring rate and adjusting the forecast accordingly.

The BIOs of *Yes-Persistence* and *Yes-Evolution* explicitly include some aspect of time in the construction of the forecasts, while the other BIOs do not add any temporal dimension. This broad distinction is the focus of testing our above-stated hypothesis.

Following Paper III, we utilize a “box-and-whisker” presentation. However, instead of focusing on skill metrics, here we focus on the frequency of occurrence of the two-day forecast outcome patterns in the context of the three two-day event histories (event/event, no-event/event, and event/no-event).

2.5. Targeted Questions

The goal of this analysis is to investigate whether/how BIOs influence the forecast outcome patterns and performance results with respect to the event histories. Specifically, we examine the following questions.

1. What is the impact of the BIOs on the independence of the two-day forecasts (meaning, does the forecast outcome for the first day significantly influence the forecast outcome for the second day)?
2. Is there any overall performance difference between BIOs within each particular categorization?
3. Do any of the BIOs better predict both the first flare and first quiet?

4. Do those BIOs that explicitly incorporate temporal information (i.e., *Yes-Persistence* and *Yes-Evolution*) display performance differences as compared to those BIOs that do not include explicit temporal information?

To address these questions, we analyze the results of the BIO performance by applying a variety of statistical methods to the forecast outcome patterns and their frequencies to answer the specific questions posed above. For example, when evaluating the influence (independence) of the first-day forecast outcome to the second-day forecast, we test the performance of the former in the context of the performance of the latter by evaluating the probability of rejecting the null hypothesis that the two are statistically independent. In contrast, when comparing the forecast performance across BIOs directly, we employ rank-sum tests, since it is solely a comparative performance that is of interest. For the four questions here, the statistical approaches are described in detail in Appendix B. The data used for this analysis and the box-and-whisker plots themselves are available (Leka & Park 2019).

3. Results

The analysis methods described above are applied to the forecasts for the participating methods, and the results are discussed below.

3.1. At-a-glance Performance Results

Forecast outcome patterns begin to emerge when the time series of forecasts is presented (Figures 3 and 4; see Figure 1).²⁰ From the *GOES* traces of the highest daily soft X-ray flux (Figures 3 and 4, top panels), it is obvious that, overall, activity is very low, as these 2 years are toward the end of the solar activity cycle. Most methods successfully predict the long intervals of no activity (i.e., gray shaded “correct nulls”), especially for the M1.0+/0/24 definition. Most methods had some intervals of correct prediction (green), and all methods had missed events (red) for both event definitions.

²⁰ Careful examination of Figures 3 and 4 reveals a slight offset in the temporal axis for SIDC and NICT, where, e.g., red shaded missed event days occasionally appear one day earlier than other methods. As mentioned in Section 2.1, custom event lists were created for these two methods due to their significantly different forecast issuance times. As discussed in Paper II, this will change the results slightly but provides our best solution to the issue at hand.

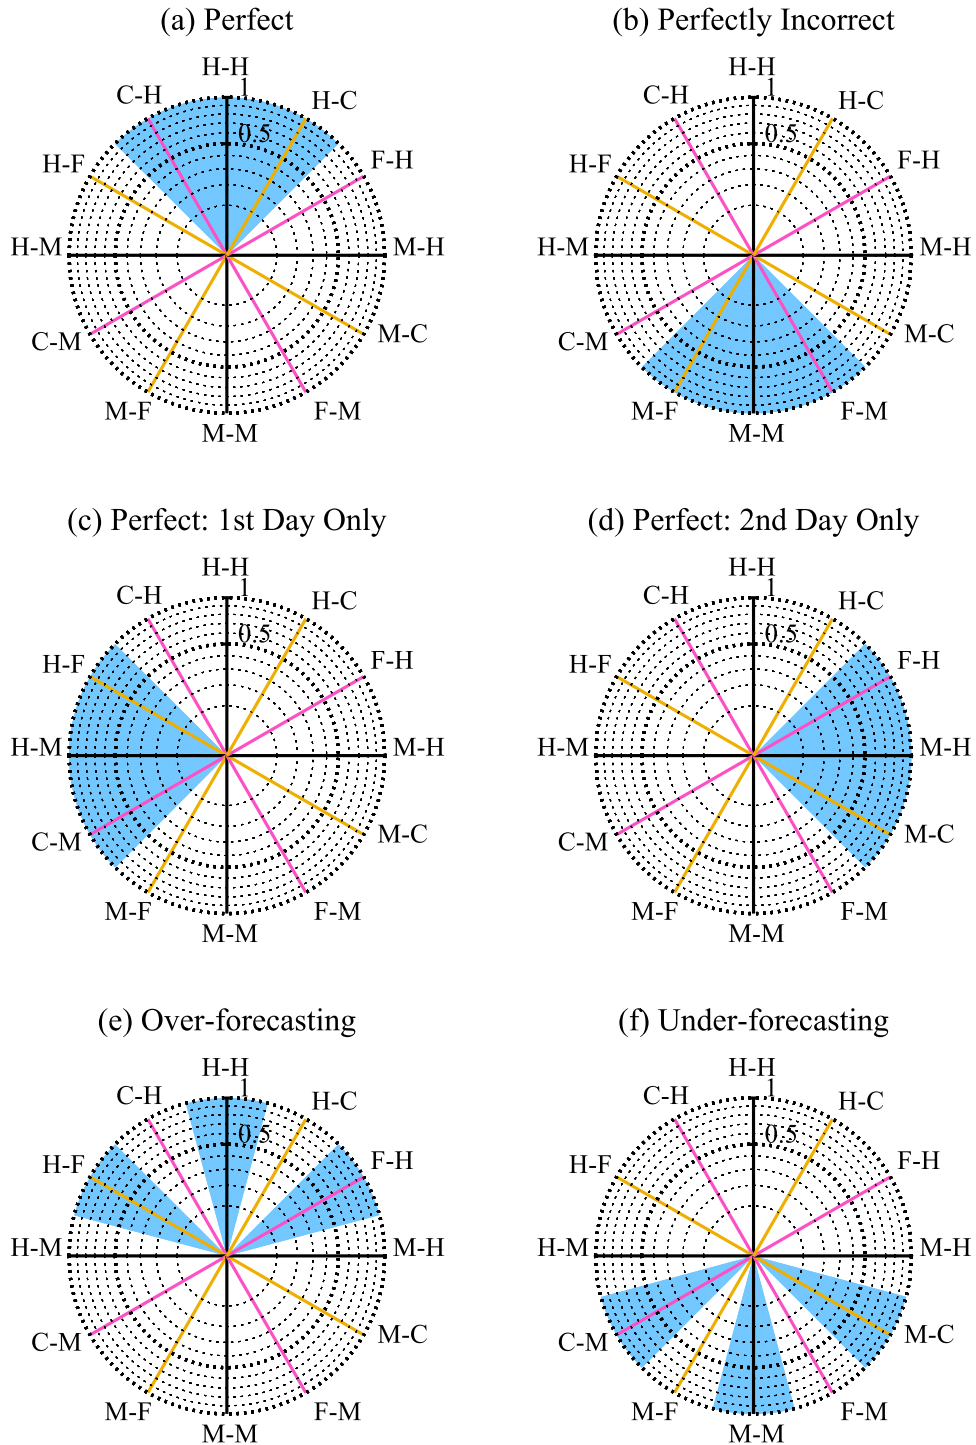


Figure 2. Radar plots are used to demonstrate the relative frequency distribution of the two-day forecasting outcome patterns. Shown are idealized cases for (a) perfect, (b) perfectly incorrect, (c) perfect only on the first day, (d) perfect only on the second day, (e) consistently overforecasting, and (f) consistently underforecasting. The two-letter axis labels indicate the direction for each of the 12 possible resulting outcome patterns (i.e., four possible combinations for each of the three event histories, as described in Table 1). The extent of the colored wedges corresponds to the magnitude of the relative frequency indicated by the concentric dotted circles at intervals of 0.1. The radius of the concentric dotted circles is determined by the square root of the relative frequency to emphasize the lower-frequency values that dominate in the present context. As an additional guide, the three event histories have color-coded axes: black (event/event), pink (no-event/event), and yellow (event/no-event). Note that the relative frequencies of the four combinations within each event-history sum to unity.

However, even with $P_{th} = 0.5$, CLIM120 (the previous 120-day-prior-climatology forecast) shows many days of false alarms (blue) for C1.0+0/24 (as do MCSTAT and MCEVOL, to a lesser degree), while NJIT shows many instances of false alarms for both event definitions. Note that CLIM120 with $P_{th} = 0.5$ was unable to make any correct M1.0+0/24

dichotomous forecasts over this testing interval; this is expected at some level, since the climatological rate is well below the P_{th} value chosen. On the other hand, for the C1.0+0/24 definition, CLIM120 produced many false alarms because the climatological rate was higher than $P_{th} = 0.5$ for the first half of the testing interval, as discussed in Paper II.

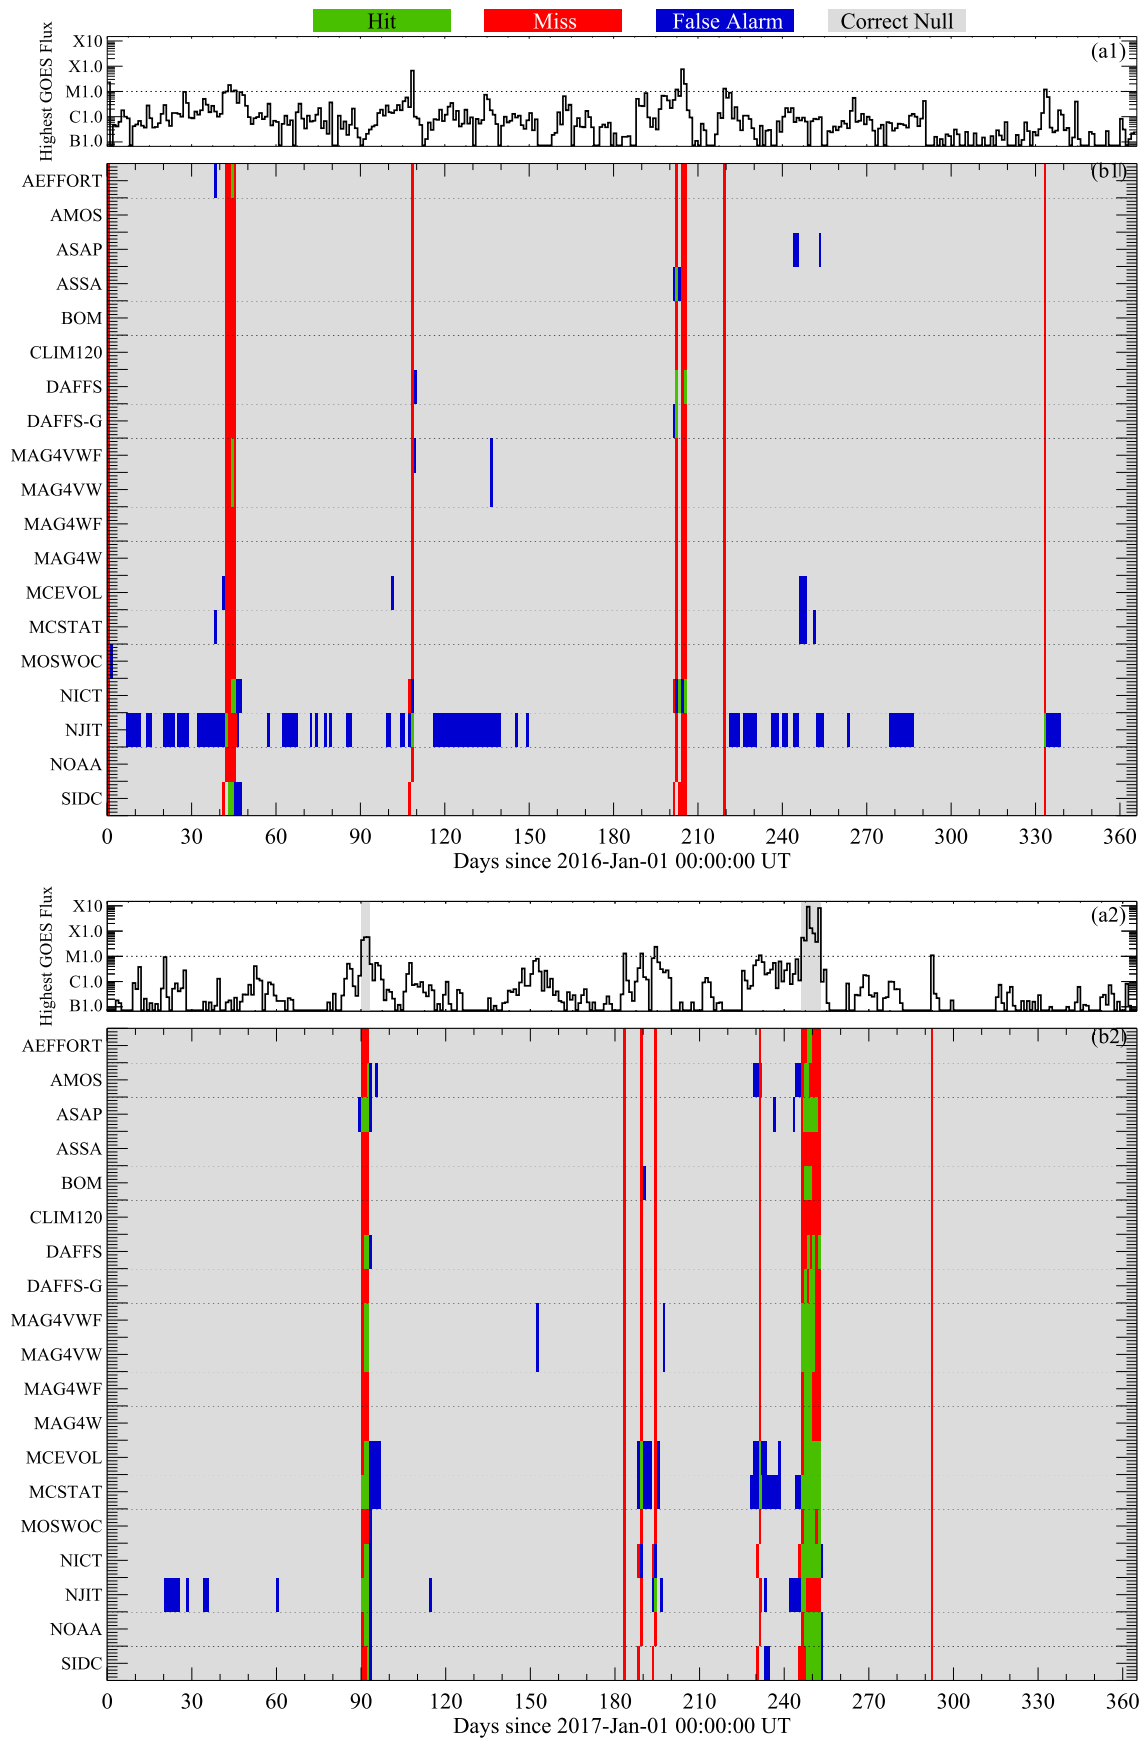


Figure 3. The daily dichotomous forecast outcomes for $M1.0 \pm 0.24$ and $P_{th} = 0.5$ are shown over the two year testing interval: 2016 (top panel) and 2017 (bottom panel). Panels (a1) and (a2) trace the highest *GOES* flux for each day, with the *GOES* M1.0 level indicated by a dotted line. Two intervals are marked with gray shading in panel (a2) 2017 April 1–3 and 2017 September 4–10, and discussed in Section 3.1. Panels (b1) and (b2) present the daily forecast outcomes—hits (green), misses (red), false alarms (blue), and correct nulls (gray)—by method, as labeled.




Figure 4. Same as Figure 3 but for daily $C1.0+0/24$ dichotomous forecasts with $P_{th} = 0.5$. The *GOES* $C1.0$ level is marked with a horizontal dotted line in panels (a1) and (a2). Note that fewer methods produce forecasts for the $C1.0+0/24$ event definition.

Two intervals are highlighted in panel (a2) of Figure 3: 2017 April 1–3 and 2017 September 4–10. These two time periods present patterns of interest for case studies due to their distinct commencement, continuation, and cessation of flaring. For the start of flaring, a no-event day is followed by an event day, and many methods correctly predict the former and miss the latter (i.e., C-M). Moving farther into the flaring interval, we have consecutive event days; for any particular event/event history, some methods miss the first but succeed or hit for the second event day (M-H), or vice versa (H-M), while others may miss both event days (M-M), which is the worst forecast outcome. Finally, approaching the end of the flaring interval, some methods correctly forecast the last event day (the “last flare”) and its following quiet day (H-C), but many follow a correct hit with a false alarm (H-F), thus not recognizing the cessation of flaring.

As discussed in Paper II, some methods lack high-probability forecasts, especially for the larger-threshold

$M1.0+0/24$ event definition. Many forecasts then register as misses for $P_{th} = 0.5$. We examine miss outcomes for larger events (i.e., $M1.0+0/24$) further in Section 3.2 in the context of $P_{th} = CLIM$ and still further in Section 3.5. Recognizing that the timelines of forecast outcomes presented here still essentially take the form of a case study (or two), we turn next to a statistical analysis of the trends.

3.2. Two-day Analysis of Forecast Outcome Patterns

As described in Section 2.3, we use a radar-plot format to analyze the performance of each forecasting method’s outcome patterns, comparing the results between $P_{th} = 0.5$ and $P_{th} = CLIM$, as well as between the two event definitions, $M1.0+0/24$ and $C1.0+0/24$, respectively, in Figures 5–8.

Referring back to Figure 2, for $M1.0+0/24$ and $P_{th} = 0.5$ (Figure 5), we see a general trend of underforecasting with a dominance of M-C and C-M outcomes (i.e., corresponding to

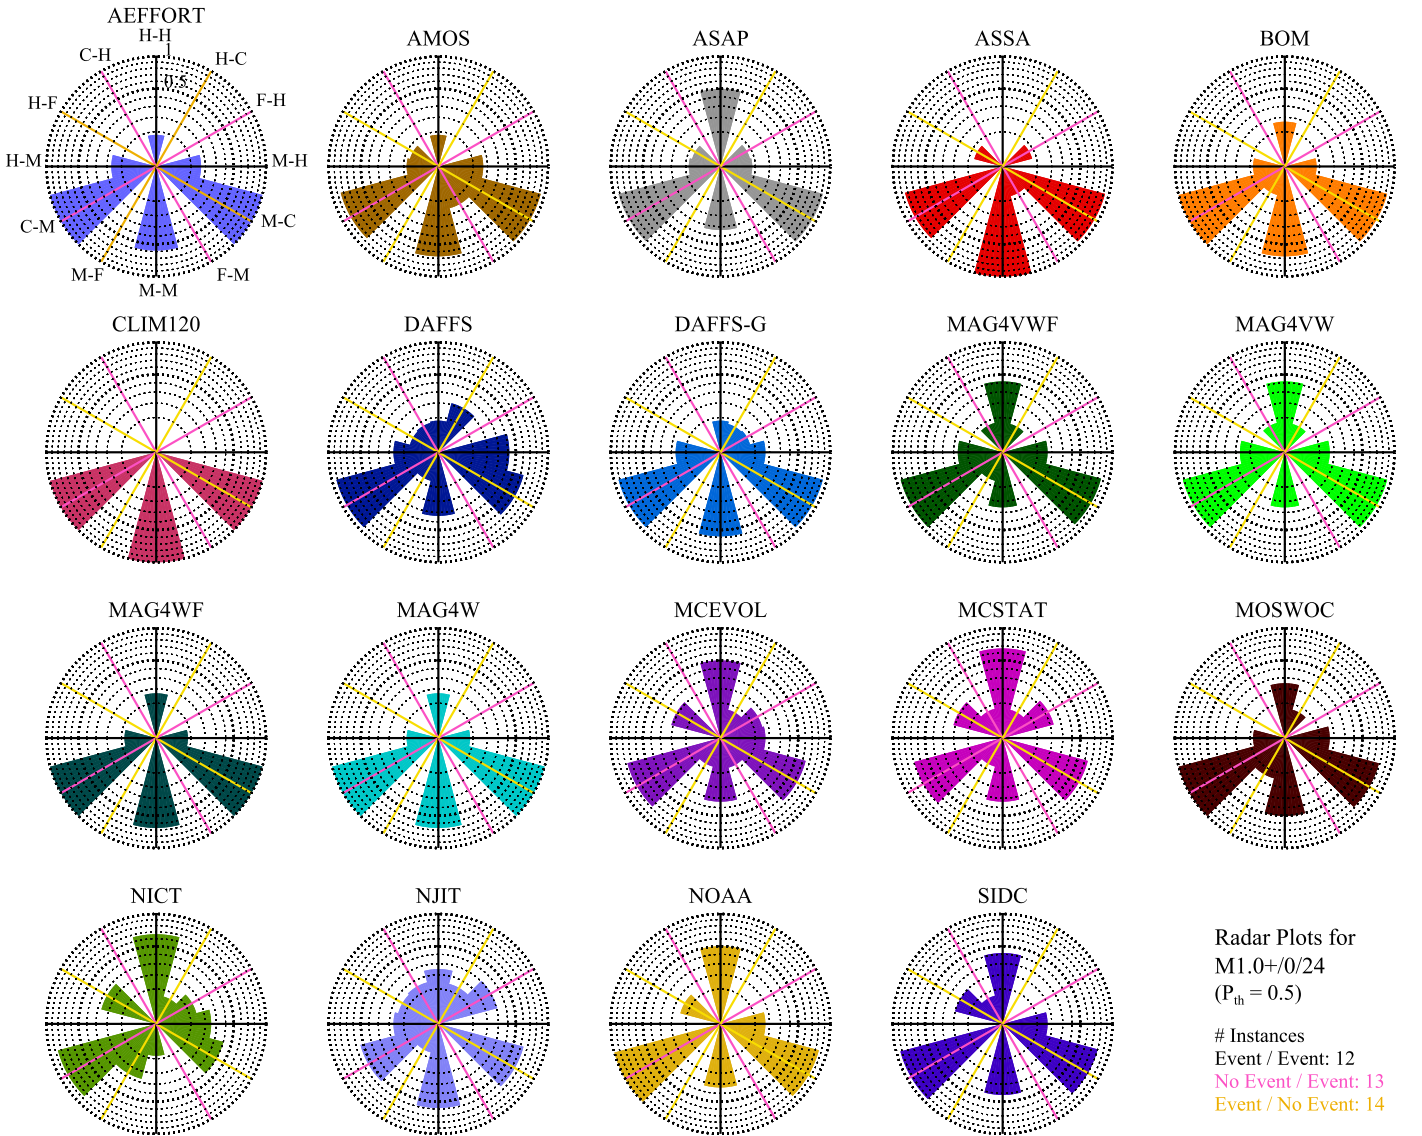


Figure 5. Radar plots for each flare forecasting method, indicating the relative frequency distribution of the two-day forecasting patterns in the M1.0+/0/24 dichotomous forecasts with $P_{th} = 0.5$ over the 2016–2017 testing interval. The colors for each method follow those used in Paper II.

04:00 and 08:00 sectors on an analog clock face, respectively) as the most frequent outcomes, with some methods additionally showing a high frequency of M-M. This analysis highlights the fact that for $P_{th} = 0.5$, most methods fail almost equally to forecast both the first flaring day of an increasing-activity period and the first flare-quiet (i.e., second) day in a decreasing-activity period. That being said, while there is a varying degree of event/event-history failure (03:00, 06:00, and 09:00 on the analog clock analogy, respectively), some methods (ASAP, MAG4VWF, MAG4VW, MCEVOL, MCSTAT, NICT, and NOAA) do show a higher H-H frequency of success than M-M frequency of failure.

The performance changes dramatically with $P_{th} = CLIM = 0.037$ (Figure 6), which is a notably low P_{th} . The trend is now overforecasting, with most methods correctly forecasting both days for the event/event history but with a high frequency of false alarms for the mixed-event histories, i.e., a dominance of F-H and H-F (02:00 and 10:00 sectors, respectively). There are now almost no complete failures for the event/event history

(i.e., a low frequency of M-M) and a few methods with almost perfect H-H results (MCSTAT, MOSWOC, and NOAA).

We next compare the radar plots for the C1.0+/0/24 forecasts with $P_{th} = 0.5$ (Figure 7) and $P_{th} = CLIM = 0.257$ (Figure 8); not all methods produce forecasts for this event definition, and following the approach in Paper II, we leave the missing radar plots blank. The threshold probability difference between 0.5 and CLIM is not as extreme as for the M1.0+/0/24 case; hence, the results are impacted less by the change in P_{th} . The same shift from under- to overforecasting is seen, especially for the mixed-event histories.

Comparing M1.0+/0/24 to C1.0+/0/24 radar plots (for those methods that produce both), the majority of methods show a higher frequency of H-H than M-M for C1.0+/0/24, which was rarely the case for M1.0+/0/24 and $P_{th} = 0.5$. In both $P_{th} = 0.5$ and $P_{th} = CLIM$ and for both C1.0+/0/24 and M1.0+/0/24, there is a consistently very low frequency for the two-day-correct outcome patterns in the mixed-event histories (C-H and H-C in the 11:00 and 13:00 sectors, respectively).

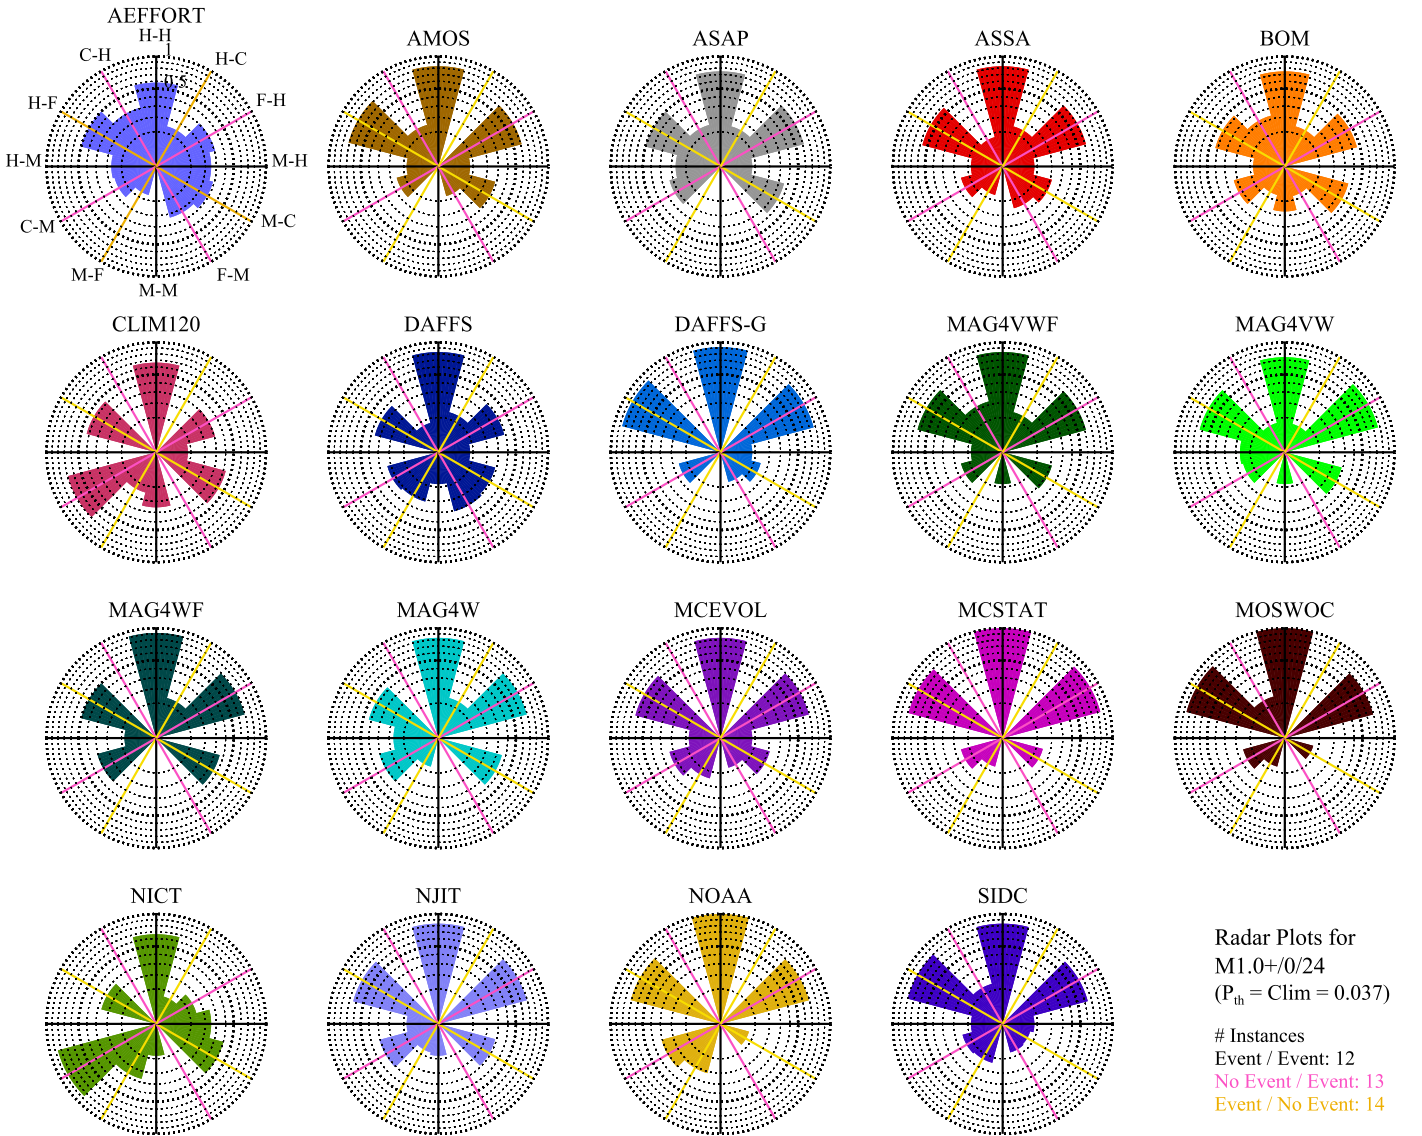


Figure 6. Same as Figure 5 but for $M1.0+0/24$ dichotomous forecasts with $P_{th} = \text{CLIM}$, where CLIM refers to the climatological rate for the testing interval (i.e., 0.036) of one or more flares occurring for the $M1.0+0/24$ event definition over the two year testing interval.

The performance of predicting the first flaring day for $M1.0+0/24$ and $P_{th} = 0.5$ is, by all accounts, poor across all methods; only NJIT even weakly registers just one of the relevant two criteria described in Section 2.3 (i.e., $C-H > F-M$). For $M1.0+0/24$ and $P_{th} = \text{CLIM}$, several methods register weak positive performance according to the criteria (i.e., $C-H > F-M$ and $F-H > C-M$). For $C1.0+0/24$ and $P_{th} = 0.5$, forecasting the first flaring day shows a modicum of success only for NJIT and only according to the second of the two criteria, i.e., $F-H > C-M$. However, for $C1.0+0/24$ and $P_{th} = \text{CLIM}$, the majority of methods can claim some success according to at least the second of the two criteria.

Forecasting the first quiet day (or the last flare) for $M1.0+0/24$ and $P_{th} = 0.5$ shows some promise for DAFFS, DAFFS-G, MAG4VM, and NJIT by the analogous criteria (i.e., $H-C > M-F$ and $M-C > H-F$), while almost all methods have some success according to one of the criteria (i.e., $M-C > H-F$) due to the prevalence of underforecasting. For $M1.0+0/24$ and $P_{th} = \text{CLIM}$, the situation changes: only NICT succeeds and only according to the second criterion, $M-C > H-F$. Turning

to the $C1.0+0/24$ definition, the results are similar to those for $M1.0+0/24$ due to underforecasting (for all methods except NJIT) at $P_{th} = 0.5$, while at $P_{th} = \text{CLIM}$, only DAFFS-G and NICT show some success according to $M-C > H-F$.

3.3. Two-day Analysis Plus Categorization: Results

The goal of the next analysis is to investigate the two-day event histories and outcome patterns in the context of differences in BIOs that summarize specific method differences suspected of influencing the results. We turn with specific interest to results from the BIOs that include explicit temporal information (i.e., *Yes-Persistence* and *Yes-Evolution*), as compared to those that include no temporal information.

The box-and-whisker plots in Figures 9 and 10 show the relative frequency distribution of the four outcome patterns (rows) for each of the three event histories (columns) for the $M1.0+0/24$ and $C1.0+0/24$ forecasts, respectively, using $P_{th} = 0.5$. We describe the results in order of fully correct (top row), fully incorrect (bottom row), and mixed errors (one of two days correct, one incorrect; middle

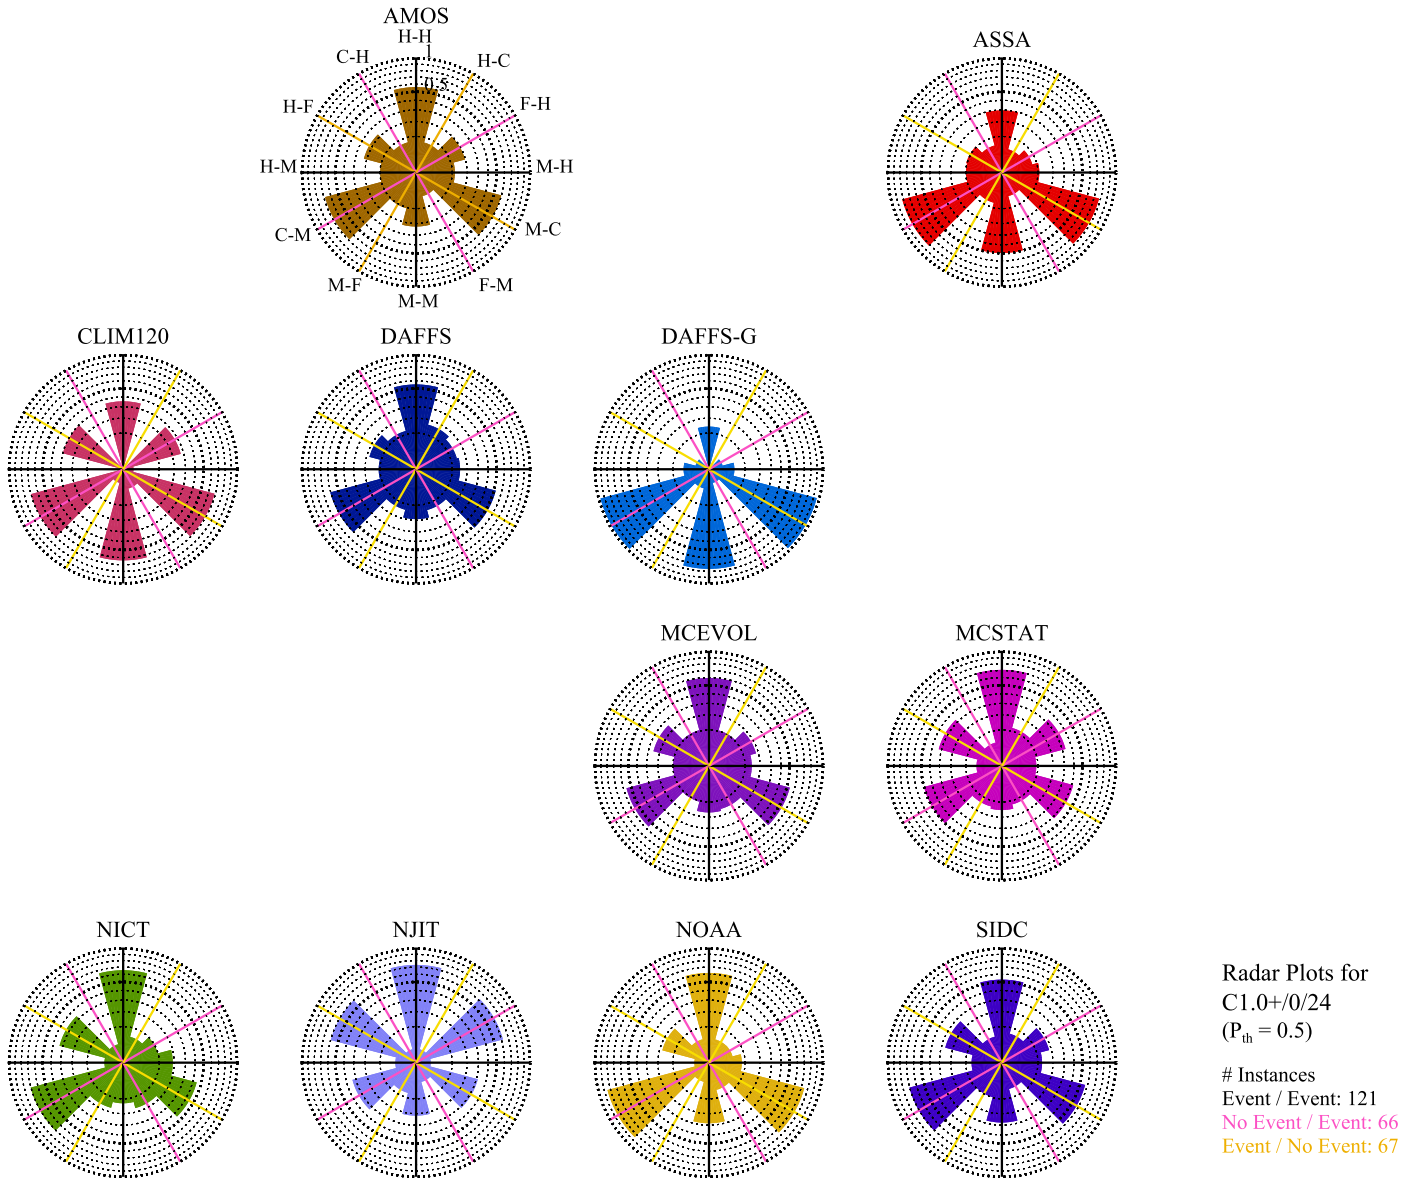


Figure 7. Same as Figure 5 but for the C1.0+0/24 dichotomous forecasts with $P_{th} = 0.5$.

two rows). Note that this order is according to outcome pattern, whereas the outcome patterns in Table 1 are listed in order according to the forecast made. The evaluations here are initially based on visual inspection of the medians of the box-and-whisker plots and secondarily based on the interquartile ranges as well. Throughout, we consider the sample size context when considering whether a result is strong or weak.

First, regarding the fully correct outcome patterns (top row of Figures 9 and 10, where a higher frequency is better), we find differences in the median values between BIOs but at more remarkable magnitudes for the event/event history than for the mixed-event histories. In the M1.0+0/24 event/event history, the BIOs of *Yes-Persistence* and *Yes-Evolution* show better performance according to the median, but the performance is similarly improved with the BIOs of *Long* and *Simple*. For the M1.0+0/24 mixed-event histories, in most cases, the medians and quartiles are very similar, with only small differences in the training and data characterization BIOs for the event/no-event history. For C1.0+0/24, which has a significantly larger event-day sample size but a smaller number of methods, the

BIOs of *Simple* and *Yes-Evolution* consistently show a slightly higher frequency (better performance) according to the medians, but this trend is diluted upon considering the quartile spreads.

Second, we turn to the perfectly incorrect outcomes (bottom row of Figures 9 and 10, where a lower frequency is better). In the case of the event/event history, there is a mirror effect as compared to the fully correct outcomes simply due to the lower frequency values of the mixed-error outcome patterns. For the M1.0+0/24 event/event history (left column), the *Yes-Persistence* and *Yes-Evolution* BIOs show a lower frequency (better performance); however, similar differences are found in the median values between the other BIOs, *Long* versus *Short/Hybrid* and *Simple* versus *Magnetic/Modern*. For the M1.0+0/24 mixed-event histories (middle and right columns), *Yes-Persistence* and *Yes-Evolution* show worse performance according to higher medians of M-F compared to *No-Persistence* and *No-Evolution*, but the same difference is seen between *Long* versus *Short/Hybrid* as well. For the C1.0+0/24 event/event history, the medians are effectively the same across all BIOs but tending toward worsening performance

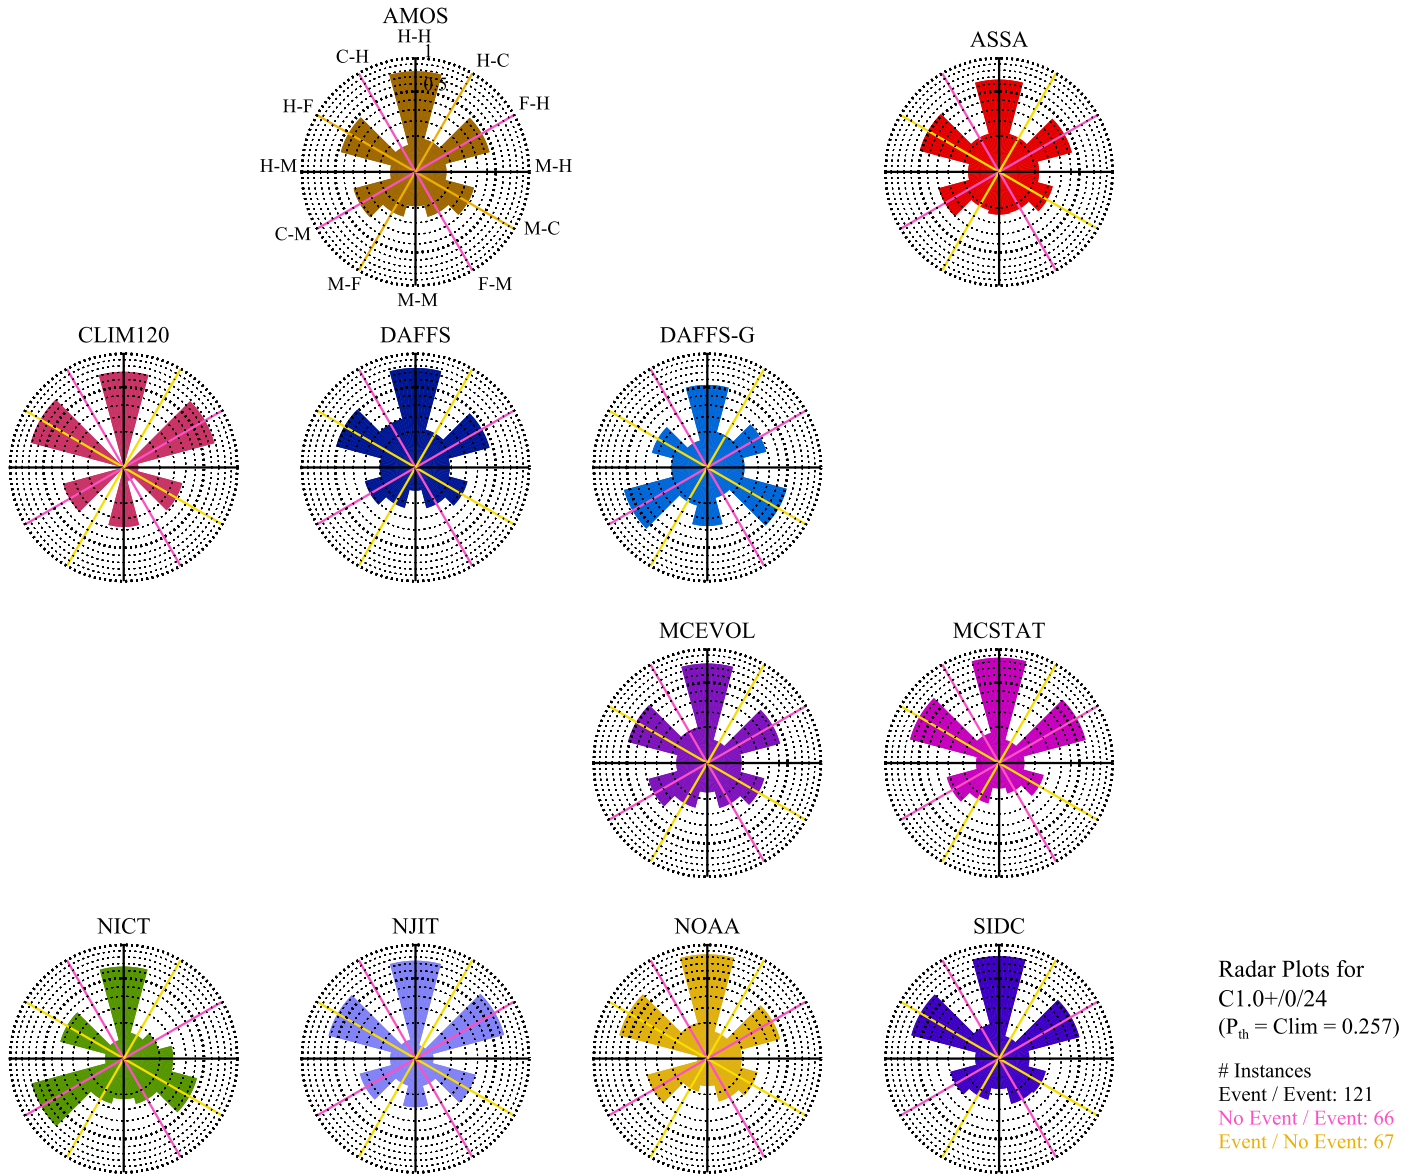


Figure 8. Same as Figure 5 but for the C1.0+/0/24 dichotomous forecasts with $P_{th} = CLIM$, where CLIM refers to the climatological event rate (i.e., 0.257) for C1.0+/0/24 over the testing interval.

(according to higher 75th percentiles of M-M) for the BIODATA that did not include temporal information (e.g., *No-Persistence* and *No-Evolution*). For the C1.0+/0/24 mixed-event histories, we may argue that the *Yes-Persistence* and *Yes-Evolution* BIODATA show a worse performance compared to other BIODATA, meaning that including temporal information leads to a false alarm after a miss or a miss after a correct null.

Finally, we summarize the mixed-error outcomes for all three event histories (middle two rows of Figures 9 and 10, where again, a lower frequency is better). Starting with M1.0+/0/24 and across the three different event histories, “first-day-correct/second-day-incorrect” outcomes (second row of Figure 9) show many significant differences between BIODATA. *Long*, *Simple*, and *Yes-Evolution* (but not *Yes-Persistence*) are advantageous for the event/event history but lead to more second-day errors for the event/no-event history. In the no-event/event history, we see a higher frequency (poorer performance) for *Yes-Persistence* and *Yes-Evolution* compared to the other BIODATA. In the case of “first-day-incorrect/second-day-correct” patterns (third row of

Figure 9), opposite trends are found in the three event histories. For C1.0+/0/24, there is a tendency that when a significant difference does occur between a pair of BIODATA for one of the mixed-event histories, the difference tends to be in the same direction between that pair for the other of the mixed-event histories. An improvement in one (e.g., C-M) is reflected in an improvement in the other (e.g., H-F). Another finding is that both *Short/Hybrid* and *Magnetic/Modern* consistently show a lower frequency (better performance) with respect to all mixed-error patterns for the mixed-event histories.

In Figures 11 and 12, the same analysis of the BIODATA is applied using $P_{th} = CLIM$. In the case of M1.0+/0/24, improvements in performance are found for the event/event history by *Long* and the event/no-event history by *Short/Hybrid*, as well as *Magnetic/Modern*. *Yes-Persistence* shows a statistically significant increase in the median (better performance) for the fully correct outcomes C-H and H-C in the mixed-event histories. On the other hand, for C1.0+/0/24, the first-day-correct/second-day-incorrect outcomes (second row) indicate that *Yes-Persistence* leads to higher

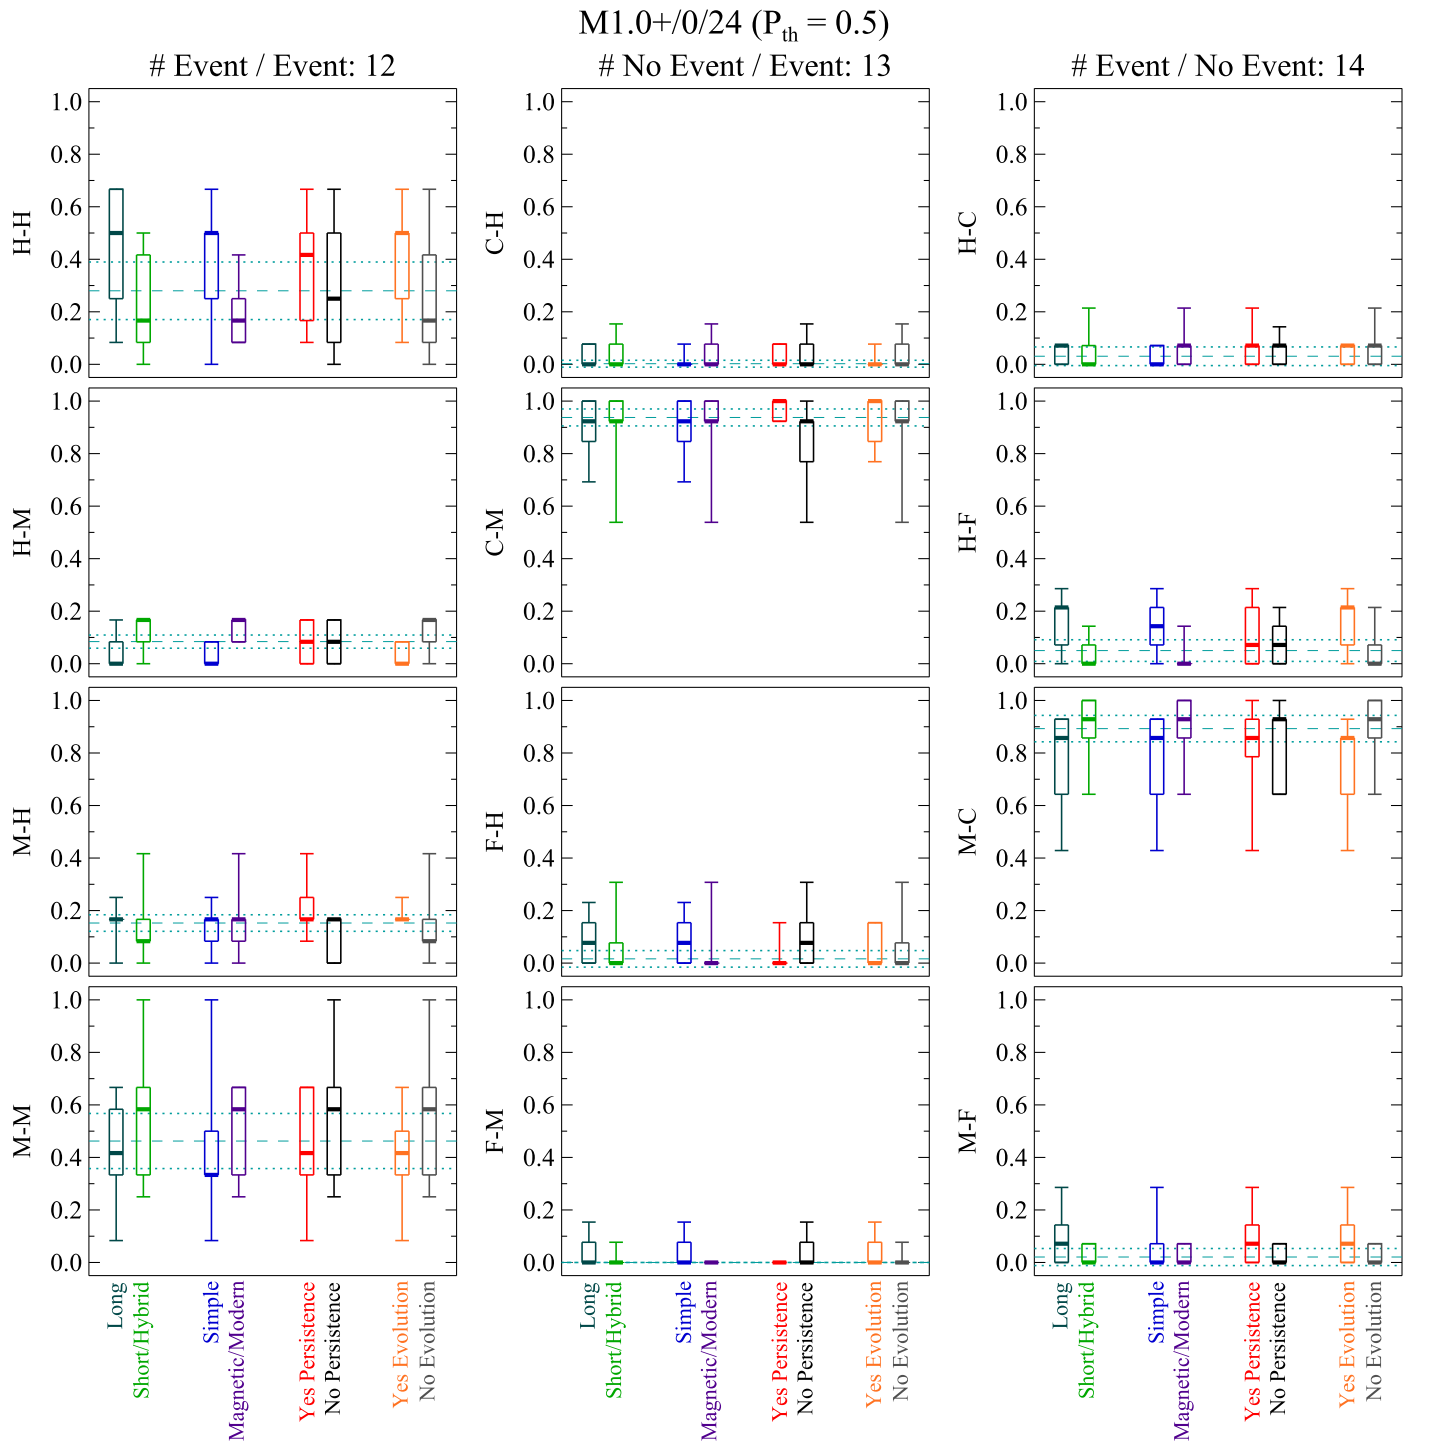


Figure 9. Comparison of the relative frequencies of the two-day forecast outcome patterns in the M1.0+0/24 dichotomous forecasts with $P_{th} = 0.5$ between different groups of methods, as categorized according to the description in the text. Box-and-whisker plots display the 25th (lower edge) and 75th (upper edge) percentiles, the median (horizontal thick line inside the box), and the minimum and maximum of the sample (whiskers). Note that if the median coincides with either the 25th or 75th percentile, that box edge will be thicker. Also plotted are the mean (dashed) \pm standard deviation (dotted) of the median values of relative frequencies from 100 sets of nine randomly selected methods among all 18 methods except CLIM120. The top row for each of the three event histories is the “correct/correct” result; hence, a higher frequency is better. All other rows indicate a frequency of results that include forecast errors; hence, lower scores are better.

errors; e.g., for H-F, acting on the first day’s activity leads to overpredicting when activity declines on the second day.

3.4. Targeted Questions Answered

We apply nonparametric statistical tests and decision trees (Appendix B) to answer the four questions presented in

Section 2.5. This approach is taken to allow more specific and quantitative analysis of the visual inspection of the box-and-whisker plots presented above.

What is the impact of the BIOs on the independence of the two-day forecasts (meaning, does the forecast outcome for the first day significantly influence the forecast outcome of the second day)? For each forecasting method, we test the null




Figure 10. Same as Figure 9 but for the C1.0+/0/24 dichotomous forecasts with $P_{th} = 0.5$. The mean (dashed) \pm standard deviation (dotted) of the median values of relative frequencies are calculated from 100 sets of five randomly selected methods among all 10 available methods for C1.0+/0/24 except CLIM120.

hypothesis, stated as, “the forecast outcome on the second day is independent of the forecast outcome of the first day.” Employing a special contingency table (see Table 8) that relates the first- versus second-day forecast outcomes, we calculate the significance level, specifically the two-sided p -value, from Fisher’s exact test (Fisher 1970) of the null hypothesis: lower p -values indicate a lower probability of accepting the null hypothesis, i.e., a higher likelihood that the day 2 forecast outcome is in fact being influenced by the outcome of the day 1 forecast. The mean of the p -values across all forecasting

methods in a given BIO is shown in each cell of Table 2 as per the event definition, P_{th} value used, and two-day event history.

The sample sizes for the M1.0+/0/24 and C1.0+/0/24 event definitions are significantly different, leading to the very different magnitudes of p -values between the two. As such, we only compare relative p -values within each definition separately and highlight relatively significant results.

With M1.0+/0/24 and $P_{th} = 0.5$, we call out *Yes-Evolution* for the event/event history and *Simple* for the event/no-event history as having the two smallest p -values but

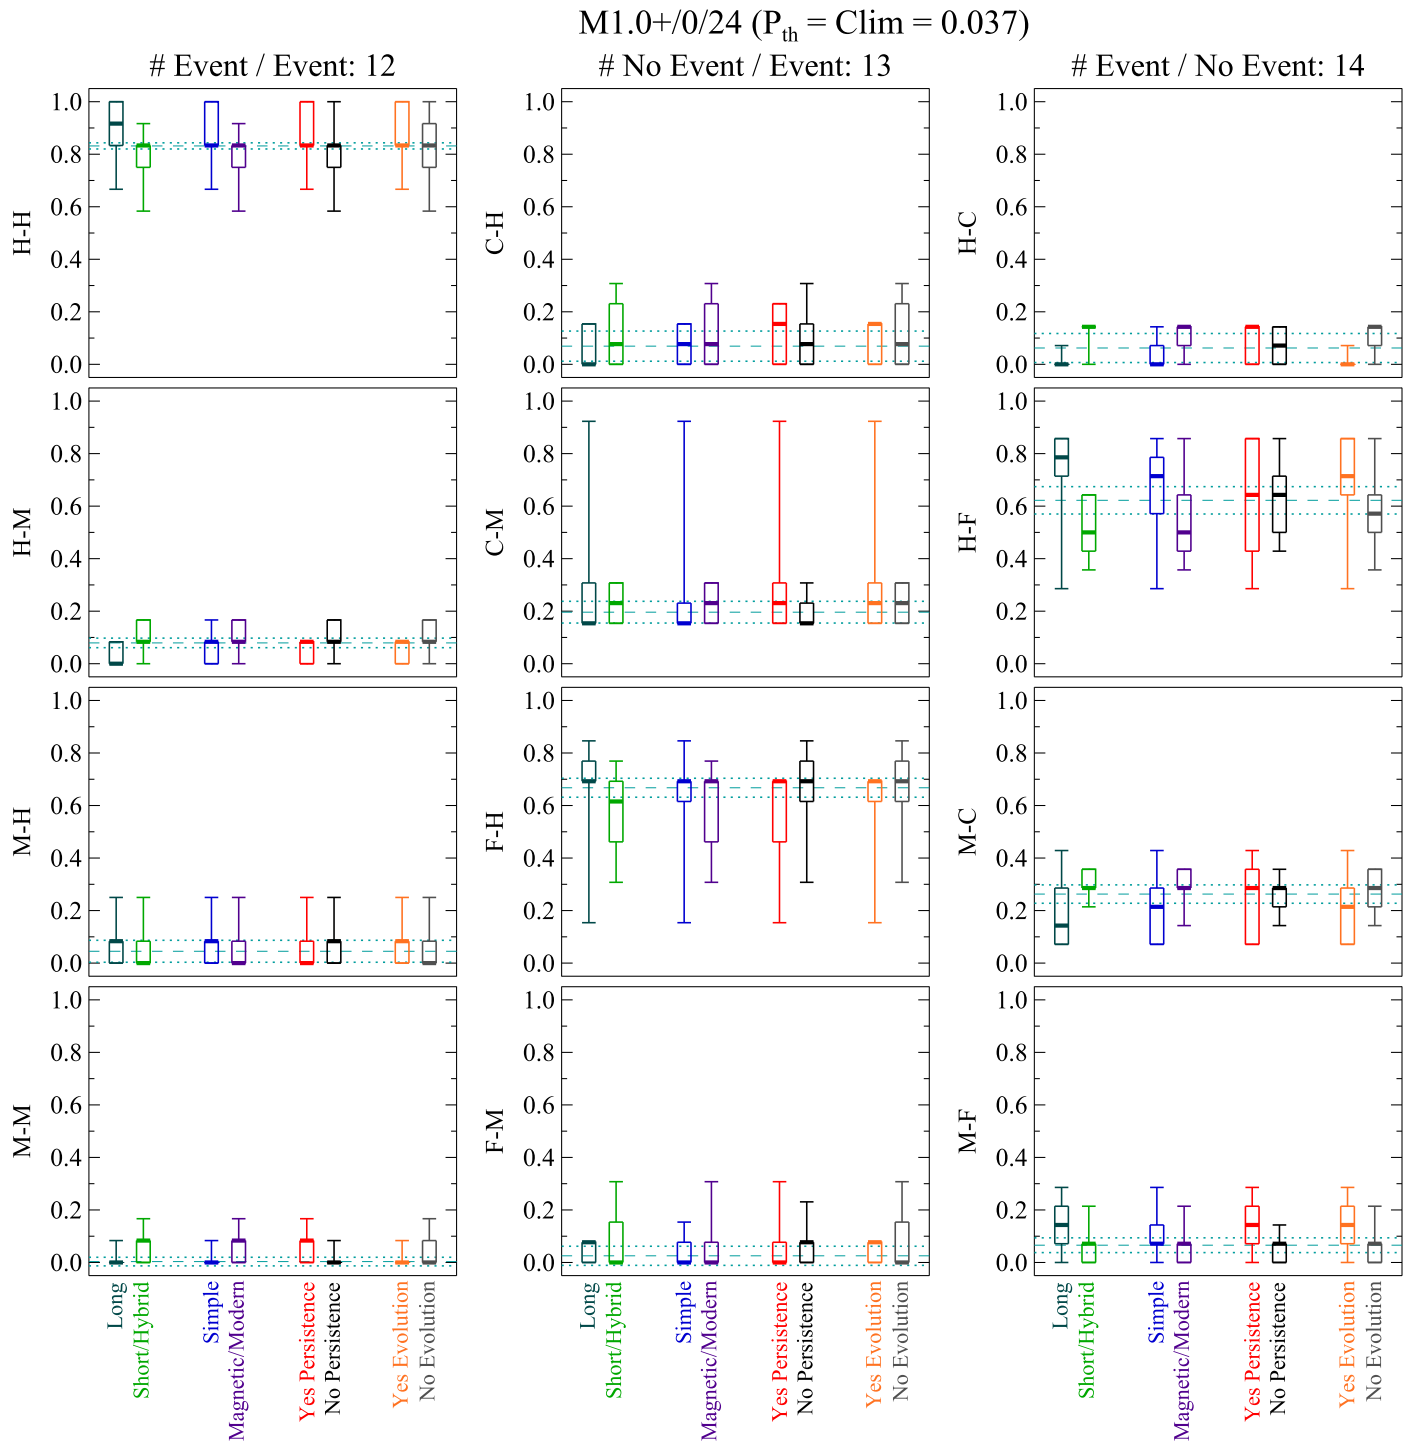


Figure 11. Same as Figure 9 but for the M1.0+0/24 dichotomous forecasts with $P_{th} = \text{CLIM}$.

not yet significant at the $p = 0.05$ level, indicating that their forecast outcomes across the two days are less likely to be independent. In Table 3, we present the contingency table entries across all methods in the BIO for these two cases, with the expected populations under the null hypothesis shown in parentheses. It is clear with this demonstration that small p -values can arise due to overpopulation of either the on-diagonal or off-diagonal elements.

For M1.0+0/24 with $P_{th} = \text{CLIM}$, large p -values are found in the event/event history across all BIOs due to the fact that the occurrence frequency of H-H is overwhelmingly higher

than that of the other three outcome patterns. This is demonstrated in the third set of entries in Table 3. On the other hand, *No-Persistence* shows a small p -value of 0.04 for the event/no-event history due to overpopulation of off-diagonal elements as shown in Table 3.

For C1.0+0/24, extremely small p -values (10^{-6} – 10^{-2}) across the BIOs result from the larger sample size coupled with either the on-diagonal or off-diagonal totals in the contingency tables always being much larger than any of the marginal totals. A tendency is found: for the event/event history, either both the day 1 and day 2 forecasts are correct or neither of them is

C1.0+/0/24 ($P_{th} = \text{Clim} = 0.257$)

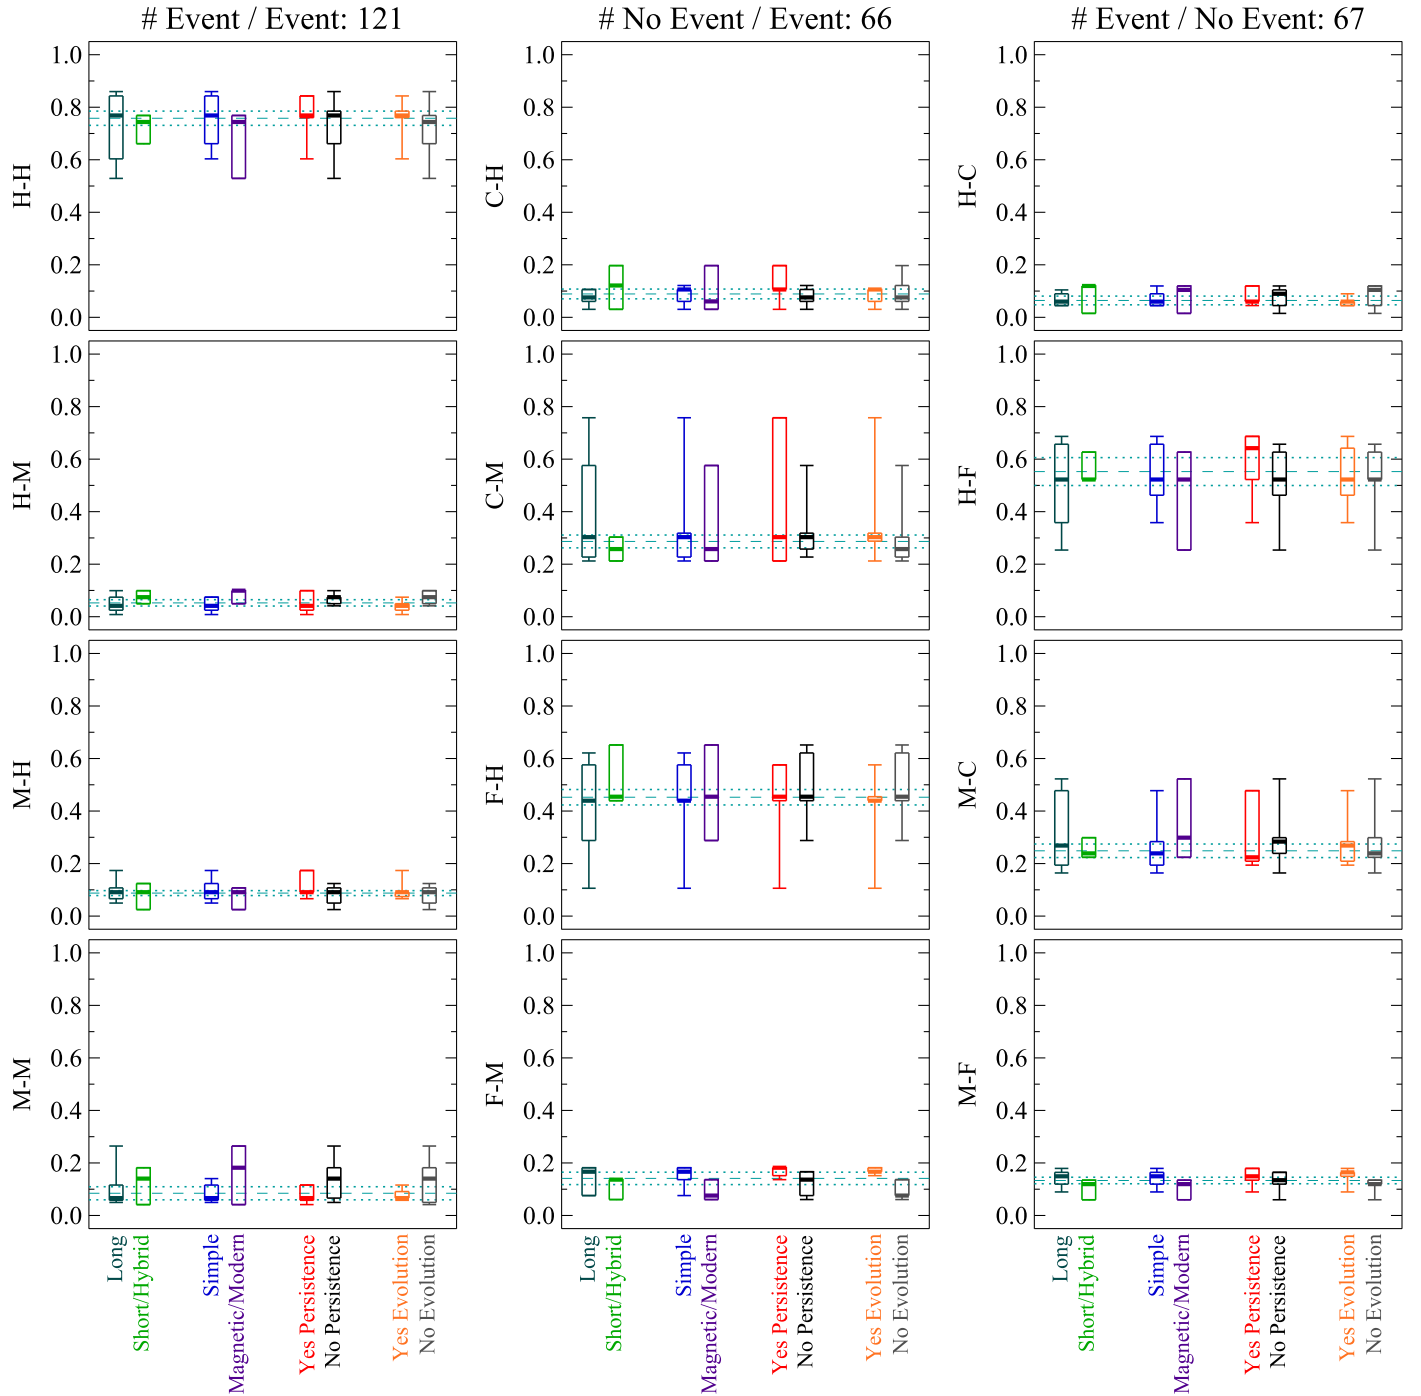


Figure 12. Same as Figure 10 but for the C1.0+/0/24 dichotomous forecasts with $P_{th} = \text{CLIM}$.

correct, while for the other mixed-event histories, only one of the two-day forecasts is correct. This leads us to further test the hypothesis that the day 1 forecast is more likely to be followed by the same day 2 forecast than if the two forecasts were independent of each other. Examining the difference between the contingency table entries and their expected values under the null hypothesis, we find across all BIOs that for the event/event history, the on-diagonal entries (H-H and M-M) consistently exceed their expected values, while for the mixed-event histories, the off-diagonal entries (F-H and C-M for the no-event/event history and M-C and H-F for the event/

no-event history) exceed their expected values. This indicates that a forecast probability higher/lower than P_{th} on day 1 tends to stay higher/lower than P_{th} on day 2 across all BIOs, as initially identified in Figure 1. It seems that forecasts do not respond fast enough (i.e., within 24 hr) to changes in the flaring history; this is a widespread failure in forecasting methods and a specific target for improvement.

Is there any overall performance difference between the BIOs within each particular categorization? This is simply “who wins?” across all outcome patterns and the implied forecasting performance between the BIOs looking across the

Table 2
Two-sided p -values from Fisher’s Exact Test of Independence for Two-day Forecasts

Event Definition	P_{th}	Event History	Training Interval		Input Parameter		Persistence		Evolution	
			Long	Short/ Hybrid	Simple	Magnetic/ Modern	Yes	No	Yes	No
M1.0+/0/24	0.5	Event/Event	0.27	0.46	0.25	0.50	0.31	0.43	0.19	0.47
		No Event/Event	0.52	0.73	0.48	0.79	0.88	0.44	0.67	0.61
		Event/No Event	0.31	0.68	0.18	0.85	0.60	0.45	0.23	0.66
C1.0+/0/24	0.5	Event/Event	2.2×10^{-6}	4.9×10^{-6}	2.2×10^{-6}	4.9×10^{-6}	4.6×10^{-6}	2.0×10^{-6}	3.1×10^{-6}	2.9×10^{-6}
		No Event/Event	6.8×10^{-3}	3.7×10^{-3}	2.6×10^{-3}	0.01	6.1×10^{-3}	5.7×10^{-3}	3.4×10^{-3}	8.3×10^{-3}
		Event/No Event	6.4×10^{-3}	4.1×10^{-3}	2.6×10^{-5}	0.02	3.1×10^{-3}	7.5×10^{-3}	3.6×10^{-6}	0.01
M1.0+/0/24	CLIM	Event/Event	0.92	0.58	0.93	0.53	0.59	0.84	0.89	0.65
		No Event/Event	0.11	0.22	0.12	0.23	0.15	0.19	0.13	0.19
		Event/No Event	0.12	0.11	0.12	0.11	0.21	0.04	0.16	0.09
C1.0+/0/24	CLIM	Event/Event	1.4×10^{-4}	0.02	1.4×10^{-4}	0.02	0.01	1.6×10^{-4}	1.8×10^{-4}	0.01
		No Event/Event	3.7×10^{-4}	6.7×10^{-3}	3.9×10^{-4}	6.7×10^{-3}	5.6×10^{-3}	9.1×10^{-5}	5.1×10^{-4}	4.0×10^{-3}
		Event/No Event	2.7×10^{-5}	2.0×10^{-4}	4.2×10^{-5}	1.7×10^{-4}	1.2×10^{-4}	5.5×10^{-5}	2.4×10^{-5}	1.4×10^{-4}

Note. The details of the calculation of the p -values are described in Appendix B.1.

event definitions and P_{th} used. The decision tree associated with this question (see Appendix B.2) is applied; a higher score indicates better performance.

As shown in Table 4, for M1.0+/0/24 with $P_{th} = 0.5$, *Long*, *Simple*, and *Yes-Persistence* (but not *Yes-Evolution*) perform slightly better, with differences in the scores against their counterpart BIOs of 0.2–0.3. For C1.0+/0/24 with $P_{th} = 0.5$, their counterpart groups (i.e., *Short/Hybrid*, *Magnetic*, *No-Persistence*, and even *No-Evolution*) perform relatively much better, with differences in the rank-based scores of 0.5–0.8.

Comparing the performance results across the event definitions and P_{th} values used, *No-Evolution* always performs better than *Yes-Evolution*, with score differences in the 0.2–1.1 range; *Yes-Persistence* performs better than *No-Persistence*, except in the case of C1.0+/0/24 with $P_{th} = 0.5$; and *Short/Hybrid* performs better than *Long*, except in the case of M1.0+/0/24 with $P_{th} = 0.5$. Relative to the maximum possible difference, those discussed here are fairly small.

Do any of the BIOs better predict both the first flare and the first quiet? This question directly addresses one of the motivations of this study, and we apply the decision tree (Appendix B.3) to achieve the results shown in Table 5, where higher totals indicate better performance. We find that *Short/Hybrid*, *Magnetic/Modern*, and *No-Evolution* attain higher total scores (2.7–3.0 out of 12), but these scores are equal to or below 25% of the maximum score possible and hence are not strong results. It is also found that most BIOs show better performance for M1.0+/0/24 than C1.0+/0/24 in the context of the first-flare/first-quiet predictions.

Do those BIOs that explicitly incorporate temporal information (i.e., Yes-Persistence and Yes-Evolution) display performance differences as compared to those BIOs that do not include explicit temporal information? To answer this final question, we reexamine Tables 2–5. We find that *Yes-Persistence* and *Yes-Evolution* show some overall performance differences compared

to the other BIOs, but the differences are not large. With respect to having a higher frequency of the two-day-correct patterns, as well as a lower frequency of the error patterns, the performance comparison between a pair of BIOs in Table 4 shows that *Yes-Persistence* performs better for M1.0+/0/24, as well as C1.0+/0/24 with $P_{th} = CLIM$, but *Yes-Evolution* performs worse across all event definitions. In addition, the performance evaluation for the mixed-event histories in Table 5 shows similar results as in Table 4. In summary, we find weak support for improvements in performance, particularly in the case of M1.0+/0/24, by explicitly including persistence or prior flare history but excluding active evolution.

3.5. Limb Events

Finally, breaking from multiday forecast outcomes, we note that the first-flare/first-quiet challenges are even more stringent when faced with very isolated flare events, i.e., when very low activity is interrupted by a single event day (see Figures 3 and 4). At the $P_{th} = 0.5$ level, there are four M1.0+/0/24 event days (of 26 event days, or 15%) for which all methods failed to provide a “yes” forecast (all methods registered a miss). These four event days share common traits: (1) only one flare event occurred on those days, (2) the source active region was located close to or behind the solar limb, and (3) solar activity was very low, with few or even no sunspots on the solar disk (see Figure 13 and Table 6). Examining these four events in some detail provides insight into the possibility of improving the forecasting in these situations.

For limb event No. 1, AR 12473 had produced M1.0+ flares a few days prior but then became quiet. At the $P_{th} = CLIM$, the majority of methods forecast M1.0+/0/24 events to occur, even though for the prior few days it had only produced low C-class events. The NJIT solely predicted a significantly higher probability (>20%) for M1.0+/0/24 when this region was at the limb (see Figure 13(a)).

Table 3
Examples of Contingency Tables for Two-day Forecasts

	Second-day Forecast	First-day Forecast	
<i>Yes-Evolution</i> (Event/Event)		Correct	Incorrect
$M1.0+0/24$ with $P_{th} = 0.5$	Correct	29 (18.1)	13 (23.9)
$p = 0.19$	Incorrect	2 (12.9)	28 (17.1)
	Second-day Forecast	First-day Forecast	
<i>Simple</i> (Event/No Event)		Correct	Incorrect
$M1.0+0/24$ with $P_{th} = 0.5$	Correct	4 (17.5)	98 (84.5)
$p = 0.18$	Incorrect	18 (4.5)	8 (21.5)
	Second-day Forecast	First-day Forecast	
<i>Long</i> (Event/Event)		Correct	Incorrect
$M1.0+0/24$ with $P_{th} = CLIM$	Correct	85 (84.3)	7 (7.7)
$p = 0.92$	Incorrect	3 (3.7)	1 (0.3)
	Second-day Forecast	First-day Forecast	
<i>No-Persistence</i> (Event/No Event)		Correct	Incorrect
$M1.0+0/24$ with $P_{th} = CLIM$	Correct	9 (31.2)	36 (13.8)
$p = 0.04$	Incorrect	88 (65.8)	7 (29.2)

In contrast, limb event No. 2 was an $M1.3$ event produced from a fast-growing active region that first appeared close to the western limb. The NOAA Edited Solar Events archive provides no information regarding location or source region for this flare. However, it was observed by both *SDO/AIA* and *PROBA2/SWAP* 174 Å (see Figure 13(b)) and was not associated with the more flare-productive ARs 12570/12572. In far-side helioseismic maps, a region can be detected at the appropriate position a few days later.²¹ The majority of methods predicted an event to occur at the $M1.0+0/24$ level for the $P_{th} = CLIM$ threshold but not at $P_{th} = 0.5$. Any significant full-disk probability was likely due to other visible regions and not from the (unassigned) source region itself.

Limb event No. 3 is very similar to event 2 in that the source region was a fast-growing emerging flux region that first appeared as it approached the western limb (see Figure 13(c)). It also appears a few days later in far-side helioseismic maps. In contrast with limb event No. 2, in this case no forecast method produced a probability of an $M1.0+0/24$ event, even at $P_{th} = CLIM$.

The last flare, limb event No. 4, occurred at the eastern limb before the source region was directly observable. The active region was not large enough to detect in the days prior using presently available far-side helioseismic maps. The NOAA Solar Region Summary Report indicated the expected return of AR 12682, which had been a fairly quiet active region on its prior disk appearance. The source region AR 12685 of limb event No. 4 produced no further flares and rapidly decayed with no other active region in the vicinity. Only one method,

the no-skill CLIM120 forecast, produced a full-disk forecast probability above the test-interval climatology.

In summary, given the small number of $M1.0+0/24$ event days, a significant fraction were missed by all methods at the $P_{th} = 0.5$ level. The majority of methods produced forecasts for an event at the $P_{th} = CLIM$ level for two of the limb flares (recognizing that $CLIM = 0.036$ is an extremely low threshold). However, all methods missed forecasting a flare day for the other two limb events, even in these almost ideal forecasting conditions (i.e., otherwise quite low activity). It is sobering to acknowledge that 15% of the event days for larger flares during this two year period were effectively beyond any forecast capability we presently have.

4. Summary and Discussion

Solar flare forecasts from a number of operational facilities worldwide have now been subjected to a set of novel evaluation methods designed to address specific behavior in the face of varying flare activity levels. The questions asked arise from the kind of information targeted in case studies: do the forecasts correctly identify a period of rising or declining activity? And if not, are there particular implementation options being used that exacerbate forecast errors of either kind (misses or false alarms) when a forecasting method is faced with temporally varying levels of flare activity? The two year testing interval targeted here (i.e., 2016–2017) includes a number of distinct periods of flare activity and quiescence, providing a good laboratory for this analysis. The performance characteristics of the forecasting methods under evaluation can be summarized as follows.

1. All methods show a trend that a high/low forecast probability on day 1 remains high/low on day 2, regardless of any observed transition between “flare-quiet” and “flare-active.”
2. Overall forecast performance is improved for $M1.0+0/24$ when persistence or prior flare history are explicitly included in computing forecasts.
3. Using magnetic/modern data leads to improvement in catching the first event day, as well as the first no-event day (for $M1.0+0/24$).
4. There are four $M1.0+0/24$ event days (of 26 event days, or 15%) during the two year testing interval for which all methods failed to provide a “yes” forecast at $P_{th} = 0.5$.

In more detail, the forecast outcomes are constructed as dichotomous forecasts by applying a threshold above which the (mostly) probabilistic output of the forecasting methods is taken to be a “yes” forecast at the given event definition (e.g., $C1.0+0/24$ or $M1.0+0/24$ in this study) or a “no” forecast when the probability values are below. We default to $P_{th} = 0.5$ as used in the earlier papers in this series, which reflects an intuitive “50/50” threshold such that a forecast probability must be above 0.5 (i.e., 50%) to be considered a forecast of an event. However, as discussed in Paper II, for larger-magnitude event definitions, most probabilistic forecast output is concentrated at low probabilities, and some methods never forecast any high probabilities. Hence, for most of the analysis methods developed herein, we also present results for $P_{th} = CLIM$ (i.e., the climatological event rate for the testing interval itself).

Because a fairly short period of two years is examined, we begin with a simple graphical depiction of the forecast outcomes in light of the full-disk soft X-ray daily maximum output (Sections 2.2 and 3.1). It is here that patterns of forecast

²¹ See http://jsoc.stanford.edu/data/farside/Composite_Maps_JPEG/.

Table 4
Performance Comparison between Each Pair of BIOs from the Same Broad Categorization

Event Definition	P_{th}	Training Interval		Input Parameter		Persistence		Evolution	
		Long	Short/ Hybrid	Simple	Magnetic/ Modern	Yes	No	Yes	No
M1.0+/0/24	0.5	2.0	1.7	2.3	2.0	1.7	1.5	1.9	2.1
C1.0+/0/24	0.5	0.9	1.4	1.1	1.8	1.1	1.9	1.0	1.7
M1.0+/0/24	CLIM	2.4	2.8	1.6	1.6	1.8	1.0	1.4	2.1
C1.0+/0/24	CLIM	1.5	2.4	2.4	2.3	2.3	1.3	1.4	2.5
Total		6.8	8.3	7.4	7.7	6.9	5.7	5.7	8.4

Note. The score of the performance comparison in each cell ranges from 0 to 12. The details of the scoring procedure are described in Appendix B.2.

Table 5
Performance Evaluation of BIOs for Mixed-event Histories

Event Definition	P_{th}	Training Interval		Input Parameter		Persistence		Evolution	
		Long	Short/ Hybrid	Simple	Magnetic/ Modern	Yes	No	Yes	No
M1.0+/0/24	0.5	0.5	0.9	0.5	1.2	0.9	0.7	0.5	1.0
C1.0+/0/24	0.5	0.5	0.3	0.5	0.3	0.5	0.4	0.5	0.3
M1.0+/0/24	CLIM	0.4	1.3	0.6	1.0	0.6	0.6	0.5	1.1
C1.0+/0/24	CLIM	0.1	0.5	0.2	0.3	0.1	0.3	0.1	0.3
Total		1.5	3.0	1.8	2.8	2.1	2.0	1.6	2.7

Note. The performance score in each cell ranges from 0 to 3; the highest achievable total score is 12. The details of the scoring procedure are described in Appendix B.3.

outcomes emerge when methods address distinct, discrete periods when solar flare activity rises, persists, and declines again over a short time period. Case studies can be useful, but any single case study may not be reflective of a method’s performance when confronted with additional such cases.

As such, we refine the analysis to focus on the critical first-flare/last-flare (i.e., first-quiet) challenge in the reality of varying flare activity. Considering two-day intervals with varying event-history cases, we specifically and statistically evaluate forecast performance in the context of increasing (no-event followed by event), continuing (event followed by event), or declining (event followed by no-event) activity histories. The radar plots presented in Section 3.2 demonstrate a quick graphical interpretation tool. From these plots, the general results of underforecasting, significant frequency of misses, and failure to correctly predict the first flare/first quiet are exquisitely clear when a probability threshold of $P_{th} = 0.5$ is applied. With $P_{th} = \text{CLIM}$, the dominant results include general overforecasting and high rates of false alarms but almost no failures on continuing-activity periods. Some methods do show an ability to correctly forecast either the first day only or the second day only, but very few methods show any ability to correctly forecast both days in the mixed-event histories (i.e., when flare activity is changing). Many methods show similar patterns to each other due to their similar approaches (previously discussed in Papers II and III).

We apply the broad categorization analysis developed in Paper III to the frequency analysis of the two-day event histories and outcome patterns in order to investigate the reasons behind certain patterns of success or failure. Box-and-whisker plots (Section 3.3) identify two-day forecasting patterns that can be interpreted in the context of the different implementations. The higher frequency of M-H when the methods include persistence or prior flare history is consistent with a forecast “adjustment”

for the second of the two days, according to the event history of the first day. This can also explain some of the two-day-incorrect patterns (i.e., F-M and M-F), as the influence of persistence or prior flare history brings the forecasts out of step with shorter timescale changes in flare activity.

We additionally ask targeted questions regarding BIOs and their effects on the independence of two-day forecasts and their performance (Section 3.4), with an emphasis on successfully predicting both the first flare and first quiet (last flare). Nonparametric statistical tests and decision-tree games were formulated to ingest the input first presented in the box-and-whisker analysis and provide quantitative answers. We first identify that for C1.0+/0/24, the day 2 forecast outcome is significantly affected by the day 1 forecast outcome across all BIOs in such a way that the two-day forecast probabilities tend to remain either higher or lower than the P_{th} for the two-day period, regardless of changes in the flaring history. We confirm weak support that including persistence or prior flare activity, as well as excluding active region evolution, improves the M1.0+/0/24 forecasts across all three two-day event histories (Table 4). On the other hand, there is evidence for improved performance of the M1.0+/0/24 forecasts when flaring activity is transitioning with the use of magnetic/modern data, even if it requires a shorter training interval (Table 5).

Except in a few cases, the results of the BIO-based analyses are not definitive. While small sample size is, of course, one culprit, another reason (as discussed in Paper III) is that the BIOs are not completely independent. As an example, the methods from RWCs (NOAA, SIDC, NICT, and MOSWOC) all employ human forecasters (i.e., FITL) and, by extension, use long training series and simple data input but also include persistence and active region evolution in their forecasts, even if in a qualitative manner. As a result, differentiation between




Figure 13. Summary images for the four at- or behind-the-limb large flares (see text). The location of each event is marked by an arrow on a full-disk composite image of the Sun obtained from the *SDO*/HMI line-of-sight magnetic field and *SDO*/AIA 131 Å (or PROBA2/SWAP 174 Å in panel (b)). The *GOES* start times and peak 1–8 Å soft X-ray fluxes are also indicated.

overlapping methods in different BIOs is diluted by the lack of true control groups where only one BIO is modified at a time.

Forecasting for flare events that occur at or behind the solar limb is known to be problematic. During the two year testing period here, four $M1.0+$ limb events were completely missed by all methods (at the 50% probability level); two of these events were correctly predicted by the majority of methods, but only at the $P_{th} = CLIM = 0.037$ level. For the other two events, all methods completely failed to produce an “event” forecast, except one instance of a correct $P_{th} = CLIM = 0.037$ event forecast from the full-disk “no-skill” 120-day-prior-climatology method. To summarize, four of 26 $M1.0+$ event days in our two year sample were missed essentially due to a lack of operationally available observations away from the Earth–Sun line.

We present here new analysis methods by which to evaluate both existing operational forecasting systems and the research and development phases of systems yet to be deployed. Specific challenges have now been presented for the flare forecasting research community beyond simply improving metrics such as those presented in Papers II and III. All operational forecasting

Table 6
Summary of Limb Flares on Four Event Days

No.	Flare		Source Region	
	Start Time	Peak Flux	NOAA Number	Location
1	2016 Jan 1 23:10 UT	M2.3	12473	S25 W82
2	2016 Aug 7 14:37 UT	M1.3	None ^a	S12 W70 ^b
3	2017 Aug 3 15:37 UT	M1.3	None ^a	N02 W85
4	2017 Oct 20 23:10 UT	M1.1	12685	S12 E88 ^b

Notes. Locations and active region assignments from the NOAA Edited Solar Events and Solar Region Summary archive (<ftp://ftp.swpc.noaa.gov/pub/warehouse>) and the *SolarSoft* Latest Events catalog (https://www.lmsal.com/solarsoft/latest_events).

^a No region number assigned before or after with which to associate this flare. See the details in Section 3.5.

^b Longitude uncertain in relevant imaging; likely behind the limb.

methods evaluated here fail to respond adequately to changes in flaring activity. As has been acknowledged in the research community (see Bloomfield et al. 2016), targeted efforts are

needed to specifically improve forecast performance over short-term variations in solar flare activity.

We wish to acknowledge funding from the Institute for Space-Earth Environmental Research, Nagoya University, for supporting the workshop and its participants. We would also like to acknowledge the “big picture” perspective brought by Dr. M. Leila Mays during her participation in the workshop. S.-H.P. gratefully acknowledges Dr. Ju Jing for maintaining the NJIT flare forecasting system and providing the archive forecasts. K. D.L. and G.B. acknowledge that the DAFFS and DAFFS-G tools were developed under NOAA SBIR contracts WC-133R-13-CN-0079 (Phase I) and WC-133R-14-CN-0103 (Phase II), with additional support from Lockheed-Martin Space Systems contract No. 4103056734 for Solar-B FPP Phase E support. A.E. McC. was supported by an Irish Research Council Government of Ireland Postgraduate Scholarship. M.K.G. acknowledges research performed under the A-EFFort project and subsequent

service implementation supported under ESA contract No. 4000111994/14/D/MPR. D.S.B. and M.K.G. were supported by the European Union Horizon 2020 research and innovation program under grant agreement No. 640216 (FLARECAST project; <http://flarecast.eu>). S.A.M. is supported by the Irish Research Council Postdoctoral Fellowship Programme and the US Air Force Office of Scientific Research award FA9550-17-1-039. The operational Space Weather services of ROB/SIDC are partially funded through the STCE, a collaborative framework funded by the Belgian Science Policy Office.

Facilities: GONG, *GOES* (XRS), PROBA2 (SWAP), *SDO* (AIA, HMI).

Appendix A Participating Methods and Facilities

In Table 7, we reproduce an abbreviated version of Table 1 from Paper II, listing the methods and facilities involved with this work and the monikers used to refer to them.

Table 7
Participating Operational Forecasting Methods (Alphabetical by Label Used)

Institution	Method/Code Name ^a	Label	Reference(s)
ESA/SSA A-EFFORT Service	Athens Effective Solar Flare Forecasting	A-EFFORT	Georgoulis & Rust (2007)
Korean Meteorological Administration & Kyung Hee University (Korea)	Automatic McIntosh-based Occurrence probability of Solar activity	AMOS	Lee et al. (2012)
University of Bradford (UK)	Automated Solar Activity Prediction	ASAP	Colak & Qahwaji (2008, 2009)
Korean Space Weather Center	Automatic Solar Synoptic Analyzer	ASSA	Hong et al. (2014), Lee et al. (2014)
Bureau of Meteorology (Australia)	FlarecastII	BOM	Steward et al. (2011, 2017)
120-day No-Skill Forecast	Constructed from NOAA event lists	CLIM120	Sharpe & Murray (2017)
NorthWest Research Associates (US)	Discriminant Analysis Flare Forecasting System	DAFFS	Leka et al. (2018)
	GONG+ <i>GOES</i> only	DAFFS-G	
	MAG4 (+according to magnetogram source and flare history inclusion)	MAG4W	
NASA/Marshall Space Flight Center (US)		MAG4WF	Falconer et al. (2011),
		MAG4VW	Appendix A in Paper II
		MAG4VWF	
Trinity College Dublin (Ireland)	SolarMonitor.org Flare Prediction System (FPS)	MCSTAT	Gallagher et al. (2002) Bloomfield et al. (2012)
	FPS with evolutionary history	MCEVOL	McCloskey et al. (2018)
Met Office (UK)	Met Office Space Weather Operations Center human-edited forecasts	MOSWOC	Murray et al. (2017)
National Institute of Information and Communications Technology (Japan)	NICT-human	NICT	Kubo et al. (2017)
New Jersey Institute of Technology (US)	NJIT-helicity	NJIT	Park et al. (2010)
NOAA/Space Weather Prediction Center (US)		NOAA	Crown (2012)
Royal Observatory of Belgium	Solar Influences Data Analysis Center human-generated	SIDC	Berghmans et al. (2005), Devos et al. (2014)

Note.

^a If applicable.

Appendix B Targeted Analysis of the BIOs

In this section, we delineate the targeted questions posed in Section 2.5 and describe the analysis method applied. For targeted questions 2 and 3 specifically, the analysis takes the form of a decision-tree “game” by which credit is applied according to a binary “win/loss” outcome, as described below.

B.1. Targeted Question 1

What is the impact of the BIOs on the independence of the two-day forecasts (meaning, does the forecast outcome for the first day significantly influence the forecast made for the second day)? To answer this, we formulate a test of the null hypothesis: the forecast outcome on day 2 is independent of the forecast outcome on day 1. In other words, given the overall frequencies of success on day 1 and day 2, does the frequency of success on day 2 depend on what was forecast and what occurred on day 1? To test this null hypothesis, we use Fisher’s exact test on a 2×2 contingency table constructed as shown in Table 8. Here a , b , c , and d correspond to the occurrence frequencies of the four outcome patterns in each event history as follows: H-H, M-H, H-M, and M-M for the event/event history; C-H, F-H, C-M, and F-M for the no-event/event history; and H-C, M-C, H-F, and M-F for the event/no-event history. This test assumes that the marginal totals (i.e., $a+b$, $c+d$, $a+c$, $b+d$) are held constant. Fisher’s exact test gives the probability of getting the observed contingency table (or a more extreme case) under the null hypothesis. A two-sided p -value is derived from this test for each forecast method as per the event definition and the two-day event history. The mean of the p -values across all methods in a given BIO is shown in each cell of Table 2.

B.2. Targeted Question 2

Is there any overall performance difference between the BIOs within each particular categorization? To answer the question, each pair of BIOs from the same categorization are compared directly. The relative performance in this context means having a higher frequency of two-day-correct patterns, as well as a lower frequency of the three error patterns. For this question, we only compare two BIOs directly and do not comment on their overall performance, only their relative performance. As such, we adopt a rank-sum approach applied to the frequency of forecast outcomes. The relative performance is then measured using the Mann–Whitney–Wilcoxon (MWW) rank-sum test, a nonparametric statistical test of the difference between two independent samples (Mann & Whitney 1947). The MWW test involves the calculation of a statistic (called U) for the two samples, respectively,

$$U_x = R_x - \frac{n_x(n_x + 1)}{2}, \quad (1)$$

where x represents the particular sample, and n_x and R_x indicate the sample size and the sum of the ranks, respectively. The absolute value of the difference between the U values for the two samples, i.e., ΔU , is then used to measure the significance of the difference between the two samples. Normalization of ΔU to a range $[0, 1]$ is achieved by dividing it by its maximum possible value for the given sample sizes (i.e., the product of the two sample sizes). Note that the maximum value of ΔU is derived from the extreme case that

Table 8
Contingency Table

		First-day Forecast	
		Correct	Incorrect
Second-day Forecast	Correct	a	b
	Incorrect	c	d

the two samples are completely separated (e.g., all values from the first sample are less than all values from the second, or vice versa). The rank-sum analysis is performed method by method on the frequency of the forecast outcomes, where the two samples are the two BIOs being compared.

The comparison and scoring are then carried out for each event definition, P_{th} value used, and BIO pair of the categories, as follows.

1. For each of the three two-day-correct patterns (i.e., H-H, C-H, and H-C), we calculate the U values for the BIO pair (e.g., A and B) and compare them (i.e., U_A and U_B).
 - (a) If $U_A > U_B$, then A will get a score of $(|U_A - U_B|)/(n_A n_B)$, while B will get no score, and vice versa. Next, go to step 2.
 - (b) If $U_A = U_B$, then both A and B will get no score. Next, go to step 2.
2. For the other nine error patterns (i.e., only first-day-correct, only second-day-correct, or all two-day-incorrect), the opposite rule is applied.
 - (a) If $U_A < U_B$, then A will get a score of $(|U_A - U_B|)/(n_A n_B)$, while B will get no score, and vice versa. Next, go to step 3.
 - (b) If $U_A = U_B$, then both A and B will get no score. Next, go to step 3.
3. Add all scores for the performance comparisons of all 12 patterns, which are shown in each cell of Table 4 as per the event definition, P_{th} value used, and BIO.

In Table 4, the score of the performance comparison in each cell ranges from 0 to 12.

B.3. Targeted Question 3

Do any of the BIOs better predict both the first flare and first quiet? This question is answered through comparisons between two different outcome patterns for the mixed-event histories only (i.e., C-H versus F-M and F-H versus C-M for the no-event/event history and H-C versus M-F and M-C versus H-F for the event/no-event history, as discussed in Section 2.3). The MWW rank-sum test is used for the comparison as described in Appendix B.2. The rules of the comparison and performance evaluation for a given BIO are as follows.

1. For the no-event/event history, the comparison of C-H versus F-M, as well as F-H versus C-M, is carried out.
 - (a) The U values (i.e., U_{C-H} and U_{F-M}) of C-H and F-M are compared.
 - i. If a given BIO has $U_{C-H} > U_{F-M}$, then it will get a score of $(|U_{C-H} - U_{F-M}|)/(n_{C-H} n_{F-M})$. Next, go to step 1b.
 - ii. If $U_{C-H} \leq U_{F-M}$, then it will get no score. Next, go to step 1b.
 - (b) U_{F-H} and U_{C-M} are compared.











- i. If $U_{F-H} > U_{C-M}$, then it will get a score of $0.5 \times (U_{F-H} - U_{C-M}) / (n_{F-H} n_{C-M})$. Next, go to step 2.
 - ii. If $U_{F-H} \leq U_{C-M}$, then it will get no score. Next, go to step 2.
2. For the event/no-event history, the comparison of H-C versus M-F, as well as M-C versus H-F, is carried out analogously.
 - (a) U_{H-C} and U_{M-F} are compared.
 - i. If $U_{H-C} > U_{M-F}$, then it will get a score of $(U_{H-C} - U_{M-F}) / (n_{H-C} n_{M-F})$. Next, go to step 2b.
 - ii. If $U_{H-C} \leq U_{M-F}$, then it will get no score. Next, go to step 2b.
 - (b) U_{M-C} and U_{H-F} are compared.
 - i. If $U_{M-C} > U_{H-F}$, then it will get a score of $0.5 \times (U_{M-C} - U_{H-F}) / (n_{M-C} n_{H-F})$.
 - ii. If $U_{M-C} \leq U_{H-F}$, then it will get no score.

In Table 5, the performance score in each cell ranges from 0 to 3.

B.4. Targeted Question 4

Do those BIOs that explicitly incorporate temporal information (i.e., Yes-Persistence and Yes-Evolution) display performance differences as compared to those BIOs that do not include explicit temporal information? There is no separate statistical test or decision tree needed to address this question. We answer by examining Tables 2–5 overall and, in particular, comparing *Yes-Persistence* and *Yes-Evolution* outcomes versus *No-Persistence* and *No-Evolution* and all other BIOs across the three prior questions.

ORCID iDs

Sung-Hong Park  <https://orcid.org/0000-0001-9149-6547>
 K. D. Leka  <https://orcid.org/0000-0003-0026-931X>
 Kanya Kusano  <https://orcid.org/0000-0002-6814-6810>
 Graham Barnes  <https://orcid.org/0000-0003-3571-8728>
 Suzy Bingham  <https://orcid.org/0000-0002-6977-0885>
 D. Shaun Bloomfield  <https://orcid.org/0000-0002-4183-9895>
 Aoife E. McCloskey  <https://orcid.org/0000-0002-4830-9352>
 Veronique Delouille  <https://orcid.org/0000-0001-5307-8045>
 Peter T. Gallagher  <https://orcid.org/0000-0001-9745-0400>
 Manolis K. Georgoulis  <https://orcid.org/0000-0001-6913-1330>

Kangjin Lee  <https://orcid.org/0000-0001-8969-9169>
 Vasily Lobzin  <https://orcid.org/0000-0001-5655-9928>
 Sophie A. Murray  <https://orcid.org/0000-0002-9378-5315>
 Rami Qahwaji  <https://orcid.org/0000-0002-8637-1130>
 R. A. Steenburgh  <https://orcid.org/0000-0001-8123-4244>
 Graham Steward  <https://orcid.org/0000-0002-9176-2697>
 Michael Terkildsen  <https://orcid.org/0000-0002-6290-158X>

References

- Barnes, G., Leka, K. D., Schrijver, C. J., et al. 2016, *ApJ*, 829, 89
 Berghmans, D., van der Linden, R. A. M., Vanlommel, P., et al. 2005, *AnGeo*, 23, 3115
 Bloomfield, D. S., Gallagher, P. T., Marquette, W. H., Milligan, R. O., & Canfield, R. C. 2016, *SoPh*, 291, 411
 Bloomfield, D. S., Higgins, P. A., McAteer, R. T. J., & Gallagher, P. T. 2012, *ApJL*, 747, L41
 Colak, T., & Qahwaji, R. 2008, *SoPh*, 248, 277
 Colak, T., & Qahwaji, R. 2009, *SpWea*, 7, S06001
 Crown, M. D. 2012, *SpWea*, 10, S06006
 Devos, A., Verbeeck, C., & Robbrecht, E. 2014, *JWSWC*, 4, A29
 Falconer, D., Barghouty, A. F., Khazanov, I., & Moore, R. 2011, *SpWea*, 9, S04003
 Fisher, R. A. 1970, *Statistical Methods for Research Workers* (14th.; Edinburgh: Oliver and Boyd)
 Gallagher, P. T., Moon, Y.-J., & Wang, H. 2002, *SoPh*, 209, 171
 Georgoulis, M. K., & Rust, D. M. 2007, *ApJL*, 661, L109
 Hong, S., Kim, J., Han, J., & Kim, Y. 2014, *AGUFM*, 2014, SH21A
 Kubo, Y., Den, M., & Ishii, M. 2017, *JWSWC*, 7, A20
 Lee, K., Moon, Y.-J., Lee, J.-Y., Lee, K.-S., & Na, H. 2012, *SoPh*, 281, 639
 Lee, S., Lee, J., & Hong, S. 2014, *ASSA GUI User Manual*, v1.10, https://www.spaceweather.rra.go.kr/images/assa/ASSA_GUI_MANUAL_110.pdf
 Leka, K. D., Barnes, G., & Wagner, E. 2018, *JWSWC*, 8, A25
 Leka, K. D., & Park, S.-H. 2019, *A Comparison of Flare Forecasting Methods II: Data and Supporting Code*, V1, Harvard Dataverse, <https://doi.org/10.7910/DVN/HYP740>
 Leka, K. D., Park, S.-H., Kusano, K., et al. 2019a, *ApJS*, 243, 36
 Leka, K. D., Park, S.-H., Kusano, K., et al. 2019b, *ApJ*, 881, 101
 Mann, H. B., & Whitney, D. R. 1947, *The Annals of Mathematical Statistics*, 18, 50
 McCloskey, A. E., Gallagher, P. T., & Bloomfield, D. S. 2016, *SoPh*, 291, 1711
 McCloskey, A. E., Gallagher, P. T., & Bloomfield, D. S. 2018, *JWSWC*, 8, A34
 Murray, S. A., Bingham, S., Sharpe, M., & Jackson, D. R. 2017, *SpWea*, 15, 577
 Park, S.-h., Chae, J., & Wang, H. 2010, *ApJ*, 718, 43
 Pesnell, W. D., Thompson, B. J., & Chamberlin, P. C. 2012, *SoPh*, 275, 3
 Sharpe, M. A., & Murray, S. A. 2017, *SpWea*, 15, 1383
 Steward, G., Lobzin, V., Cairns, I. H., Li, B., & Neudegg, D. 2017, *SpWea*, 15, 1151
 Steward, G. A., Lobzin, V. V., Wilkinson, P. J., Cairns, I. H., & Robinson, P. A. 2011, *SpWea*, 9, S11004
 Woodcock, F. 1976, *MWRv*, 104, 1209