Published in final edited form as:

Phys Rev E Stat Nonlin Soft Matter Phys. 2014 August ; 90(2): . doi:10.1103/PhysRevE.90.022133.

Statistical evaluation of forecasts

Malenka Mader^{1,2,3,4,5,*}, **Wolfgang Mader**^{2,3,4}, **Bruce J. Gluckman**^{4,5,6,7}, **Jens Timmer**^{2,3,8},
and **Björn Schelter**^{2,9,†}

¹Department of Neuropediatrics and Muscular Disease, University Medical Center of Freiburg, Germany

²Freiburg Center for Data Analysis and Modeling (FDM), University of Freiburg, Germany

³Institute for Physics, University of Freiburg, Germany

⁴Center for Neural Engineering, Pennsylvania State University, PA, State College, USA

⁵Department of Engineering Science and Mechanics, Pennsylvania State University, PA, State College, USA

⁶Department of Neurosurgery, Pennsylvania State University, PA, Hershey, USA

⁷Department of Biomedical Engineering, Pennsylvania State University, PA, State College, USA

⁸BIOSS, Center for Biological Signalling Studies, University of Freiburg, Germany

⁹Institute for Complex Systems and Mathematical Biology, University of Aberdeen, UK

Abstract

Reliable forecasts of extreme, but rare events, such as earthquakes, financial crashes, or epileptic seizures, would render interventions and precautions possible. Therefore, forecasting methods have been developed which intend to raise an alarm if an extreme event is about to occur. In order to statistically validate the performance of a prediction system, it needs to be compared to the performance of a predictor randomly raising alarms. Such a random predictor can be obtained by bootstrapping or analytically. We propose an analytic statistical framework which, in contrast to conventional methods, allows for validating independently sensitivity and specificity of a forecasting method. Moreover, our method accounts for the periods, during which an event has to remain absent or occur after a respective forecast.

I. INTRODUCTION

Extreme but rare events occur in ecology as in absence/presence models of species [1], in weather as in form of storms [2, 3], heavy precipitation [2, 3], floods [4], and extreme temperatures [3], in geology as earthquakes, and in economics [5] as financial crises, as well as in medicine in form of strokes or epileptic seizures [6, 7].

*Corresponding author: M.Mader@fdm.uni-freiburg.de. †Second corresponding author: B.Schelter@abdn.ac.uk.

PACS numbers: 02.50.-r, 95.75.Wx

The impacts of such extreme and unpredicted events are severe. A reliable forecasting system could alleviate those impacts, because precautions or interventions [6, 7] could be taken. A forecasting system would be considered reliable when predicting both an upcoming event (sensitivity) and the absence of a future event (specificity). Depending on the application, the focus might vary on either a high sensitivity or high specificity of the prediction.

In general, a forecasting system is based on several features, which quantify different properties of observations, such as frequency content, regularity, etc., for reviews, see [6, 8, 9]. These features are supposed to quantify a change of the observed system before an extreme event occurs. In that case a probabilistic [2, 10, 11] or a binary forecast [1, 3, 6, 9, 12] can be derived from such features.

For probabilistic forecasts, the probability of the occurrence of an event is provided [2, 10, 11]. Such forecasts leave the user with the decision about the necessity to initiate precautions or preventions [5, 13]. Probabilistic forecasts can be statistically assessed by the Brier Score [14]

$$B = \frac{1}{N} \sum_{t=1}^N (\hat{p}(t) - I(t))^2, \quad (1)$$

which is the temporal mean of the squared deviations of the probabilities $\hat{p}(t)$ of an upcoming event at time t to the indicator function $I(t)$ of actually occurring events. The indicator function $I(t)$ is zero, except at times of events when $I(t) = 1$. Inserting a binary forecast

$$\hat{I}(t) = \begin{cases} 1 & \text{if an alarm is raised at time } t \\ 0 & \text{else} \end{cases} \quad (2)$$

instead of the probabilistic forecast $\hat{p}(t)$ into Eq. (1), yields

$$B = \frac{1}{N} \left(\sum_{\tilde{t} \in \text{TP}} 0 + \sum_{\tilde{t} \in \text{FP}} 1 + \sum_{\tilde{t}'' \in \text{TN}} 0 + \sum_{\tilde{t}'' \in \text{FN}} 1 \right) \quad (3)$$

$$= \frac{m_+ \cdot 0 + (T - m_+) \cdot 1 + m_- \cdot 0 + (N - T - m_-) \cdot 1}{N} \quad (4)$$

$$= 1 - \frac{m_+ + m_-}{N}, \quad (5)$$

where $m_{\pm}$ are the numbers of true positives and true negatives, T is the total number of predicted events, and N is the number of prediction-observation pairs available. For this, the sum in Eq. (1) is split up into four sums, as in Eq. (3). These sums contain the time points of true positive (TP), false positive (FP), true negative (TN), and false negative (FN) predictions $\hat{I}$.

Equation (5) shows that for the Brier Score the number of true positives m_+ and true negatives m_- are weighted equally, when applied to binary forecasts. This is particularly inexact when faced with rare events because a high number of true negatives then obscures a low number of true positives.

For binary forecasts, contingency tables are typically considered. They contain information about true and false positives, as well as true and false negatives. The statistics of the contingency table can be assessed by the Fisher–Irwin test [15]. For the evaluation of a prediction, a time frame within which the event must or must not occur needs to be defined. This time frame enhances the probability of true positive and true negative predictions, which is not accounted for in the statistics of the Fisher–Irwin or similar tests. As an alternative to statistically testing the entries of contingency tables directly, certain measures [12, 16, 17] and according statistical tests [18, 19] are available that take into account the time frames in which events must or must not occur. However, they do not quantify the fraction of true positives, i.e., the sensitivity, and the fraction of true negatives, i.e., specificity, independently [12, 16, 17]. The statistics [19] of the seizure prediction characteristics [18], for instance, is based on a Poissonian random predictor with binomial success distribution when assuming a certain false prediction rate, i.e. specificity. This dependency of sensitivity and specificity is also shown for the Brier Score in Eq. (5). To summarize, neither for binary nor for probabilistic forecasts independent statistical assessment of sensitivity and specificity in the framework of predictions has been proposed so far. We close this gap by presenting an adaptive Poissonian random predictor to obtain binomial success distributions for the sensitivity and specificity, independently.

The article is organized as follows. In Sec. II we review the statistical assessment of binary forecasting methods. In Sec. III we propose and test a statistical method to independently assess the number of true positives and true negatives of forecasting methods based on random predictors. For two different settings we show the adaptability of our method, Secs. IIIA and III B. By simulations we test the efficiency of the proposed statistics for both settings. We complete this paper by discussing the method proposed (Sec. IV).

II. STATISTICAL PROPERTIES OF BINARY PREDICTION METHODS

A binary forecast, which we denote a prediction in the following, can either be positive or negative. A positive prediction yields an alarm raised to announce that an event is about to occur. If correctly predicting, a positive prediction leads to a true positive. True negatives are

negative predictions followed by the absence of events. Standard ways of quantitatively assessing true positives and false positives, true negatives and false negatives are described in Sec. II A.

In Sec. II B, the idea of statistical evaluation of the contingency tables, is reviewed. The performance of the prediction system needs to be tested in terms of superiority to a random predictor. We propose an analytic solution for this step.

A. Quantification of Prediction-Observation Matches

A useful framework for quantifying the matches of predictions and observations, is based on the introduction of an occurrence and intervention period (Fig. 1). The period during which the anticipated event has to occur after an alarm is the occurrence period (OP) [6]. This renders a classification of true positives (Fig. 1, a) and false positives (Fig. 1, b) possible. Typically, an event-free period is desirable after the alarm before the OP sets on. It allows for an intervention or precaution to be initiated after the alarm, before the occurrence of the event (Fig. 1 green, dotted box). This period is denoted intervention period (IP) [6, 8], prediction horizon [20], or lead-time [4].

A negative prediction (Fig. 2) is followed by the IP in the sense that a non-occurrence of an event should be guaranteed during this phase to obtain a true negative prediction (Fig. 2, a). If an event occurred on the other hand, the negative prediction would be wrong, leading to a false negative (Fig. 2, b).

The respective numbers of true positives and false positives, true negatives and false negatives are summarized in contingency tables, which then can be statistically validated, as described in the next section.

B. Evaluation of Prediction-Observation Matches

To statistically validate contingency tables derived from the assessment of positive and negative predictions (Sec. II A), chi-square tests or the Fisher–Irwin test can be used to test independence of predictions and events [15, 21]. However, they do not account for the statistical effect of the OP. After a positive prediction an event can occur within the duration of the OP. For a long OP, the probability of a true positive prediction is increased compared to a short one. In order to circumvent this statistical challenge, measures [1, 6, 9, 12] based on the contingency tables have been introduced. They quantify the performance of the predictor by weighting sensitivity and specificity, differently. One example of such a measure is the fraction correct, which is the sum of sensitivity and specificity [9]. For some of those measures, the distributions of the test statistics are known, for others they are approximated by Gaussian distributions. Those distributions of test statistics are obtained for the case of linkage and the case of independence of prediction-observation pairs. In case of linkage, the predictions of a prediction method are correlated to the observations. On the other hand predictions and observations are uncorrelated in case of independence. Correct predictions might still occur, for instance a positive prediction is followed by an event. However, such correct predictions are due to chance. The aim of any statistical assessment of prediction methods is to exclude that the performance is just as good as it is for such an independent prediction method, for which predictions and observations coincide by chance.

That is, the null hypothesis of independence of predictions and observations needs to be rejected. To this end, a threshold which separates the two distributions corresponding to linkage and independence for a given error probability, can be obtained analytically [9]. Since error probabilities of positives and negatives are linked by the choice of the threshold, an independent statistical assessment of sensitivity and specificity is not possible. In contrast to the methods used so far, we propose an analytic method of validating predictions with respect to the temporal extent during which predicted events can occur. Moreover, the proposed method renders independent statistical control of true positives and true negatives possible.

III. STATISTICAL ASSESSMENT OF PREDICTIONS WITH INDEPENDENT CONTROL OF SENSITIVITY AND SPECIFICITY

In this section, we propose a statistics for the evaluation of the performance of a prediction method based on the number of true positives and true negatives with respect to actually observed events.

The statistics proposed is based on the idea to evaluate the performance of a prediction method, by comparing it to the performance of a random predictor. Here, we assume a homogeneous Poissonian random predictor that raises alarm at a constant rate but independent of the actually occurring events. Adaptions to other types of random predictors are necessary if other null-hypotheses shall be tested. In case, e.g., it is known that the events occur at a time-dependent rate such that the numbers of events vary throughout day, month or year, the random predictor needs to fulfill the same time-dependence. Then the null-hypothesis of performing quite as good as the prediction method investigated given that time-dependence is tested. Such adaptions, however, do not change the idea of the subsequent statistics. For binary random predictions, the number of correct predictions, i.e., accordances of predictions and observations, are binomially distributed. Therefore, a threshold of the maximum number of correct predictions obtainable by chance can be determined from a binomial distribution. This leads to a binomial statistics, which can be adapted to different settings of testing the accordance of predictions and observations.

To assess the performance of the statistics proposed, prediction-observation pairs are simulated. These prediction-observation pairs mimic the series of prediction-observation pairs derived from any prediction method for which the performance should be evaluated. Since in the simulations the linkage of predictions and observation are known, the effectiveness of the statistics is investigated.

In the following, two prediction scenarios are investigated, illustrating the adaptivity of the binomial distribution ansatz. In the first scenario, a prediction method with an OP of one sample point and without IP is simulated, Sec. III A. In the second scenario, IPs and OPs are simulated as extended periods in time, Sec. III B. For each scenario, the two steps of proposing the statistics and assessing its effectiveness are conducted.

A. Evaluating Prediction-Observation Pairs without IP and short OP

A prediction method performs significantly better than chance if it exceeds the performance of a random predictor given a certain error probability. The random predictor raises alarms randomly throughout time, independent from properties used in the investigated prediction method to predict the observed events. Given the alarm rate is the same as the alarm rate of the prediction method investigated, the probability of a positive prediction before an event is

$$p_+ = \frac{n_a}{N}, \quad (6)$$

the fraction of number of alarms n_a of the prediction method and the number of data points N of the prediction-observation time series. Both properties can be derived from the time series of predictions and observations.

Since the random predictor can be either right or wrong, the number of true positive predictions is binomially distributed. The probability for m_+ true positive predictions out of M_+ positive predictions is thus

$$\lambda_+ = \mathcal{B}(m_+, M_+; p_+) = \binom{M_+}{m_+} p_+^{m_+} (1-p_+)^{M_+-m_+}. \quad (7)$$

From this binomial distribution the quantiles are derived from the inverse of the cumulative binomial distribution Λ^{-1} . This leads to the critical value

$$m_{c+} = \Lambda^{-1}(1-\alpha_+, M_+; p_+), \quad (8)$$

which is the maximum number of true positives, obtainable by the random predictor with an error probability α_+ . This error probability is the significance level.

The critical value m_{c-} of true negatives is obtained analogously, with the number of negative predictions $M_- = N - M_+$ and a negative-prediction probability,

$$p_- = (1-p_a), \quad (9)$$

instead of M_+ and p_+ in Eq. (8), and p_a the probability for an alarm. If desired, a different significance level α_- instead of α_+ can be used for the control of the number of true negatives.

In order to assess the statistical properties of the investigated prediction method the following four-step algorithm has to be applied.

1. Determine the significance level $\alpha_{\pm}$ for true positives and true negatives, respectively. This is the error probability of deciding that the prediction method is better at positive and negative predictions than chance, although it actually is not.
2. Estimate the probability for a positive and negative prediction $p_{\pm}$ based on the number of positive and negative predictions, n_a and $N - n_a$ in the prediction time series.
3. Derive the critical values of true positives m_{c+} and true negatives m_{c-} from the binomial distributions, Eq. (8), with significance levels $\alpha_{\pm}$, numbers of positive and negative prediction $M_{\pm}$, and positive and negative prediction probabilities $p_{\pm}$, Eqs. (6) and (9).
4. Compare the critical values $m_{c\pm}$ obtained in step 3. to the number of true positives m_+ and negatives m_- obtained from the prediction-observation time series as described in Sec. II A. If the number of true positives and true negatives exceeds the respective critical values the prediction system performs significantly better than chance.

In the following subsections we present a simulation study. First the simulation of data is presented. Second the simulation study is described and results are presented.

1. Simulation of Data—In order to test the statistics proposed in Sec. III A, we simulated prediction-observation time series as in [10]. To this end, first a time series of probabilistic forecasts f was drawn from a β -distribution with density

$$s(f) = \frac{f^{v-1}(1-f)^{w-1}}{F(v, w)}, \quad (10)$$

where $F(v, w)$ is the β -function and v, w parameters of the distribution. From the parameters the expected value of probabilistic forecasts can be derived, $\mu_f = \frac{v}{v+w}$. The β -distribution is constrained to the range $[0, 1]$, and thus can be thresholded to zeros and ones. To obtain binary predictions a ($a = 1$ alarm vs. $a = 0$ no alarm), the forecasts f were thresholded such that the probability for an alarm was μ_f . To this end, the threshold was selected to be the inverse of the cumulative β -distribution with parameters v and w evaluated at $1 - \mu_f$. This ensured the expected value of the forecast to be the same as for the prediction, such that $\mu_f = \mu_a$.

In the second step, the probability for an event, i.e., $x = 1$, after a prediction a was modeled by

$$\mu_{x|a} = ca + d \quad (11)$$

a linear linkage c of predictions a to the probability of an event after a prediction $\mu_{x|a}$ with an offset $d = \mu_x(1 - c)$ (Fig. 3), as in [10]. The offset d was chosen such that the expected value of an event was equal to the expected value of predictions, $\mu_x = \mu_a$. Thus, depending on whether an alarm was raised ($a = 1$, Fig. 3 red, solid line) or not ($a = 0$, Fig. 3 blue, dashed line), the probability of an event $\mu_{x|a}$ was a linear function of the linkage strength c . The observation x was obtained based on the prediction a by drawing a random number x' of a uniform distribution on the interval $[0, 1]$ and comparing it to $\mu_{x|a}$ such that

$$x = \begin{cases} 1 & \text{if } x' < \mu_{x|a} \text{ (event occurs)} \\ 0 & \text{else (no event occurs)} \end{cases}. \quad (12)$$

This way, prediction-observation pairs with known linkage could be generated. At high linkage the probability of an event after a positive prediction (red, solid line) was high, while the probability of an event after negative prediction (blue, dashed line) was low. At low linkage strengths ($c \approx 0$) the probability of events was similar no matter whether a positive or negative prediction was made. While zero-linkage thus refers to independence of predictions and observations, such that the null-hypothesis is met, increasing linkage refers to increasing probability of accordances of predictions and observations.

2. Simulation Study—To assess the performance of the proposed statistics, we simulated prediction-observation time series consisting of $N = 10000$ prediction-observation pairs as described in Sec. IIIA1. Investigating particularly the case of few events, in which measures like the Brier Score fail, we chose the probability of an event to be $\mu_x = \mu_a = 0.05$ with $v = 1$ and $w = 19$ (Sec. IIIA1). Linkage strengths were increased from no linkage, $c = 0$, to maximum linkage, $c = 1$. That way, it was known whether the null-hypothesis or the alternative actually applied to the data. This yielded power analysis possible. Power refers to the probability of the statistics to reject the null-hypothesis when actually the alternative is true. Applied to the prediction-observation pairs assessed here, the power is expected to increase with increasing linkage strength. On the other hand, coverage was investigated. Coverage refers to the probability to reject the null-hypothesis, when it actually is true. To ensure that coverage is kept correct, the probability of erroneously rejecting the null-hypothesis should not exceed the significance level. Coverage and power of the statistics lead to a reliable and powerful statistics.

For such power and coverage analysis, we simulated $K = 1000$ repetitions of prediction-observation time series, each $N = 10000$ data points long, for a set of linkage strengths $c \in [0, 1]$. The first 500 of 10000 data points of exemplary prediction-observation time series are shown in the inset of Fig. 4 for various linkage strengths (y-axis). Positive predictions, i.e. alarms, are shown as red long vertical lines, while positive observations, i.e. events, are shown as blue short lines. While at zero-linkage alarms and events occur independently, the number of overlapping alarms and events increase with increasing linkage strength. According to our four-step algorithm we statistically assessed the performance of the simulated prediction-observation pairs. The probability of the rejection of the null-

hypothesis of independence of prediction-observation pairs was then derived from the percentage of repetitions for which the null-hypothesis of this independence was rejected.

In Fig. 4 the results of the power analysis both for true positives (red, solid line) and true negatives (blue, dashed line) are shown for linkage strengths $c \in [0, 1]$. For absent coupling $c = 0$, the number of true positives and true negatives remains below the significance level $\alpha_{\pm} = 0.05$ (black, dotted line). Therefore, the statistics is considered reliable. Since for true negatives, the percentage is below the significance level, the proposed statistics is conservative. With increasing linkage strength, the power of the statistical test proposed increases quickly, both for true positives and true negatives. This yields a powerful statistics.

As the simulations show, the statistics of the number of true predictions based on binomial distribution is reliable and powerful in assessing the performance of a prediction method.

B. Evaluating Prediction-Observation Pairs with extended IP and OP

If the predictions are evaluated using the framework of temporally extended IPs and OPs, their durations T_o and T_i need to be incorporated into the statistics.

The critical values for true positives (+) and true negatives (−) are derived from the inverse of a binomial cumulative distribution, as before, Eq. (8), with different prediction-observation probabilities p_+ and p_- . As shown in Fig. 1, for a true positive prediction, no event must occur in the IP but at least one event has to occur in the OP. This yields a probability

$$p_+ = (1 - p_a)^{T_i} \cdot (1 - (1 - p_a)^{T_o}) \quad (13)$$

of a true positive prediction. After a true positive prediction, the next prediction can be made after the end of the event.

Similar to the true positive prediction, the probability of a true negative prediction can be derived. For a true negative prediction (Fig. 2), the occurrence of no event must be ensured within the IP. Therefore, given a negative observation, the probability of T_i negative predictions is $p_- = (1 - p_a)^{T_i}$. However, at the next time point the probability of a true negative is increased since it is already known from the previous time point that $T_i - 1$ of the next predictions are negative. Thus the probability of a true negative prediction at the next time point is

$$p_- = (1 - p_a). \quad (14)$$

In the case of few events, this is an approximation of a global probability of true negative predictions.

The probabilities of true positives p_+ and true negatives p_- describe the probability of the random predictor to predict the observations, correctly. From the inverse binomial cumulative distribution, $\Lambda^{-1}(1 - \alpha_{\pm}, M_{\pm}; p_{\pm})$, the critical values of true positives and true negatives can be derived as in Sec. III A. The four-step algorithm presented in Sec. III A, thus can be applied when substituting $M_{\pm}$ and $p_{\pm}$ by the values obtained when incorporated extended durations of IP and OP as presented here.

1. Simulation of Data—Similar to Sec. IIIA1, we simulated prediction-observation time series based on β -distributed forecasts. In contrast to the simulations above, the observations were constructed from the thresholded forecasts such that the probability of an event was low during the IP, and high in the OP after a positive prediction. Additionally, we introduced a refractory period (RP), which occurs after each event. This simulates the period after the event, during which no predictions can be made because e.g. the system investigated does not provide reliable predictions during or in a transition time after the event since the measurements are still affected.

The parameters of the β -distribution were changed to ensure that the probability of an event during a period as long as the IP-OP complex was lower than 10%. This yielded parameters $\nu = 1$, $w = 199$, and $\mu_f = 0.005$. Since now the OP was extended, the number of events was decreased under the null-hypothesis of independent predictions and observations. To ensure $\mu_a = \mu_x$ as before, events were added randomly at instances of no events, whenever the number of alarms exceeded the number of events. This kept the null-hypothesis correct, i.e., if no linkage was present. For increasing linkage, the number of events converges to the number of alarms due to linkage. In order to leave the exact number of added events flexible, we drew uniformly distributed random numbers between 0 and 1 for all time points at which no event occurred. The random numbers were thresholded at $\frac{\Delta}{\# \text{ no events}}$ such that the expected number of additional random events is the difference of number of alarms and events. As in Fig. 4, the first 500 of 10000 data points of exemplary prediction-observation time series are shown for various linkage strengths in the inlet of Fig. 5. As before, the correlation of alarms (red long) and events (blue short vertical lines) increases with increasing linkage. As expected by the simulation design, the alarms tend to precede the events in contrast to the simulation corresponding to the inlet in Fig. 4. Furthermore, since the alarm rate is reduced in this simulation, fewer alarms and events occur.

2. Simulation Study—As described in Sec. IIIA2 we performed power analysis based on 1000 repetitions of prediction-observation time series including IPs and OPs as described in Sec. IIIB1. We assumed the IP to be shorter than the OP, choosing $T_i = 2$ and $T_o = 4$ data points. In our simulations, we chose the RP to be as long as the IP, such that $T_r = 2$. The resulting power curve for linkages $c \in [0, 1]$ is shown in Fig. 5.

Both true positives (red, solid line) and negatives (blue, dashed line) kept the coverage at $c = 0$ correct. Again, the statistics is conservative for true negatives, since the percentage of rejections are lower than 5%. For increasing linkage strength, the power increases quickly both for true positives and true negatives. Thus, the statistics could be adapted to the case of extended IPs and OPs, remaining powerful while keeping coverage correct.

IV. DISCUSSION

We introduced a statistical method, that evaluates the performance of binary forecasting systems by a random predictor. Other than previous statistical methods, it allows independent evaluation of true positive and true negative predictions, even when extended phases such as the occurrence and intervention periods are used.

We proposed an algorithm for the statistical assessment of predictions obtained from a prediction method. This algorithm includes the proposed statistics of a random predictor. The random predictor is built such that the properties of prediction equal those of the investigated prediction method except for the linkage of prediction and observations. For this random predictor the statistics of obtaining true positive and true negative predictions by chance is given by a binomial distribution. From this distribution the quantiles of the random predictor can be derived and compared to the number of true positives and true negatives given by the prediction method. Here, we described how critical values of true positives and true negatives can be obtained from the respective inverse of the cumulative binomial distribution. If the number of true predictions obtained from the prediction method exceeds the number of true predictions of the random predictor, the prediction method is considered significantly better than chance. Instead of computing critical values, p -values can be obtained from the binomial cumulative distribution. Instead of computing the inverse of the cumulative binomial distribution evaluated at 1 minus the significance level α , the cumulative binomial distribution has to be evaluated at the number of true predictions made by the prediction method investigated.

In two sections, we showed that our statistics is adaptive to frameworks based on occurrence and intervention periods. Furthermore, we showed, that the statistics proposed is powerful by testing the performance in controlled settings. Power is high to identify prediction methods which are linked to the observations such that predictions are better than chance. On the other hand, coverage is kept correct in cases, in which the prediction method does not predict observations better than chance.

Our framework proposed here also implies the assessment of strengths of events, e.g., strengths of thunderstorms or earthquakes. To this end, the event strengths can be categorized into several groups. For each group the performance of true positives and true negatives then can be assessed with the statistics proposed. Alternatively, the binomial statistics can be adapted to multinomial instead of binomial distributions.

Furthermore the binomial statistics proposed can be applied to the statistical assessment of event-detection. In that case, intervention and occurrence period can be replaced by respective periods of interest. Adaptions of the statistics are then necessary. They are implementable analogously to the ones presented here.

As the simulations show, the proposed method is a reliable, powerful, adaptive and widely applicable statistics for assessing the performance of prediction or detection methods, even in case of rare events.

Acknowledgments

This research was supported in part by The US National Institute of Health through grant R01NS065096, as well as E!5185 PANG.

References

1. Fielding AH, Bell JF. Environmental Conservation. 1997; 24:38.
2. Murphy AH. Weather and Forecasting. 1991; 6:302.
3. Frei C, Schär C. Journal of Climate. 2001; 14:1568.
4. Lardet P, Obled C. Journal of Hydrology. 1994; 162:391.
5. Jolliffe, IT., Stephenson, DB. Forecast verification: A practitioner's guide in atmospheric science. 2. Wiley; West Sussex, UK: 2012.
6. Mormann F, Andrzejak RG, Elger CE, Lehnertz K. Brain. 2007; 130:314. [PubMed: 17008335]
7. Schelter, B., Timmer, J., Schulze-Bonhage, A. Seizure Prediction in Epilepsy: From Basic Mechanisms to Clinical Applications. Wiley; Weinheim, Germany: 2008. p. 334Chap. Seizure Prediction in Epilepsy: Does a Combination of Methods Help?
8. Feldwisch-Drentrup H, Schelter B, Jachan M, Nawrath J, Timmer J, Schulze-Bonhage A. Epilepsia. 2010; 51:1598. [PubMed: 20067499]
9. Marzban C. Weather and Forecasting. 1998; 13:753.
10. Bradley AA, Schwartz SS, Hashino T. Weather and Forecasting. 2008; 23:992.
11. Jachan, M., Feldwisch-Drentrup, H., Posdziech, F., Brandt, A., Altenmüller, D-M., Schulze-Bonhage, A., Timmer, J., Schelter, B. 4th European Conference of the International Federation for Medical and Biological Engineering. In: Sloten, JV,Verdonck, P,Nyssen, M., Hauelsen, J., editors. IFMBE Proceedings. Vol. 22. Springer; Berlin Heidelberg, Germany: 2009. p. 1701-1705.
12. Doswell CA III, Davies-Jones R, Keller DL. Weather and Forecasting. 1990; 5:576.
13. Casati B, Wilson LJ, Stephenson DB, Nurmi P, Ghelli A, Pocernich M, Damrath U, Ebert EE, Brown BG, Mason S. Meteorological Applications. 2008; 15:3.
14. Brier GW. Monthly Weather Review. 1950; 78:1.
15. Fisher RA. Journal of the Royal Statistical Society. 1922; 85:87.
16. Marzban C. Journal of Applied Meteorology. 1998; 37:72.
17. Wilks, DS. International Geophysics Series: Statistical Methods in the Atmospheric Sciences. 2. Dmowska, R,Hartmann, D., Rossby, HT., editors. Vol. 91. Academic Press Inc, Elsevier Inc; Burlington, USA: 2005. p. 648
18. Maiwald T, Winterhalder M, Aschenbrenner-Scheibe R, Voss HU, Schulze-Bonhage A, Timmer J. Physica D. 2004; 194:357.
19. Schelter B, Winterhalder M, Maiwald T, Brandt A, Schad A, Schulze-Bonhage A, Timmer J. Chaos. 2006; 16:013108. [PubMed: 16599739]
20. Winterhalder M, Maiwald T, Voss HU, Aschenbrenner-Scheibe R, Timmer J, Schulze-Bonhage A. Epilepsy and Behavior. 2003; 4:318. [PubMed: 12791335]
21. Campbell I. Statistics in Medicine. 2007; 26:3661. [PubMed: 17315184]

FIG. 1.

A prediction method that raises an alarm (red, bold vertical line) can be validated only if the event has to occur in a specific time window (blue, solid box), the occurrence period (OP) of length T_o , after the alarm was triggered. The period between the alarm and the onset of the OP is the intervention period (IP) of length T_i in which intervention systems can be initiated (green, dotted box). A true positive alarm (a) is then defined as an alarm which is followed by an event in the OP. If an event remains absent (b) in the OP the alarm is a false positive.

FIG. 2.

Negative predictions refer to the intervention period (IP, green, dotted box), since the time points after the intervention period could belong to an occurrence period of a positive prediction after the negative prediction. A true negative prediction (a) refers to a negative prediction after which no event occurs within the intervention period. A false negative prediction (b) refers to one after which an event occurs at the end of the intervention period.

FIG. 3.

The prediction method simulated here is based on a linear function of the linkage c (x-axis) of predictions a (red, solid or blue dashed line) to the probability of an event $x = 1$ (y-axis). While the probability of an event increases linearly with the linkage strength for positive predictions ($a = 1$, red, solid line), it decreases for negative ones ($a = 0$, blue, dashed line), see Eq. 11.

FIG. 4.

Power curve of the simulated prediction system when compared to the proposed random predictor. Power was tested on 1000 repetitions for each of linkage strengths $c \in [0, 1]$, for both true positives (red, solid line) and true negatives (blue, dashed line). The significance level is shown in black (dotted line).

FIG. 5.

Power curve of the simulated prediction system when compared to the proposed random predictor. Power was tested with 1000 repetitions for linkage strengths c , for both true positives (red, solid line) and true negatives (blue, dashed line) when intervention period and occurrence period were used.