Published in final edited form as:
J Clin Neurophysiol. 2013 April ; 30(2): 174–177. doi:10.1097/WNP.0b013e3182767c35.

Inferring Seizure Frequency From Brief EEG Recordings

M. Brandon Westover[1,*], **Matt T. Bianchi**[2], **Mouhsin Shafi**[1], **Daniel B. Hoch**[1], **Andrew J. Cole**[1], **Keith Chiappa**[1,3], and **Sydney S. Cash**[1]

[1]MGH Epilepsy Service, Massachusetts General Hospital, and Harvard Medical School, Boston, MA 02114

[2]Department of Neurology, Sleep Division, Massachusetts General Hospital, and Harvard Medical School, Boston, MA 02114

[3]King Edward VII Memorial Hospital, Bermuda

Abstract

Routine EEGs remain a cornerstone test in caring for people with epilepsy. Although rare, a self-limited seizure (clinical or electrographic only) may be observed during such brief EEGs. The implications of observing a seizure in this situation, especially with respect to inferring the underlying seizure frequency, are unclear. The issue is complicated by the inaccuracy of patient-reported estimations of seizure frequency. The treating clinician is often left to wonder whether the single seizure indicates very frequent seizures, or if it is of lesser significance. We applied standard concepts of probabilistic inference to a simple model of seizure incidence to provide some guidance for clinicians facing this situation. Our analysis establishes upper and lower bounds on the seizure rate implied by observing a single seizure during routine EEG. Not surprisingly, with additional information regarding the expected seizure rate, these bounds can be further constrained. This framework should aid the clinician in applying a more principled approach toward decision making in the setting of a single seizure on a routine EEG.

Keywords

Bayes'rule; seizure frequency; statistical inference; EEG

INTRODUCTION

Most epilepsy centers rely on brief (<30 minutes) electroencephalogram (EEG) recordings for diagnostic purposes. Occasionally a patient has a seizure during the study. Such an event might suggest a high underlying seizure rate, because otherwise the chance of observing a seizure during any brief interval is low. Alternatively, one might argue that a single seizure may have simply been a rare event in a patient with a low underlying seizure rate, and that occasional seizures are expected in a high-volume EEG service by chance alone. An

[*]Corresponding author, Massachusetts General Hospital, 55 Fruit Street, WACC 827, Boston, Massachusetts 02114, mwestover@partners.org, Phone: 617-726-3311, Fax: 617-726-9250.

DISCLOSURE OF CONFLICTS OF INTEREST

None of the authors has any conflict of interest to disclose. We confirm that we have read the Journal's position on issues involved in ethical publication and affirm that this report is consistent with those guidelines.

accurate estimation of the underlying seizure frequency has important clinical implications, dictating whether to pursue further testing, recommend medication changes, or continue a watchful waiting approach. We explore these opposing perspectives within a principled probabilistic framework, and show that each may be appropriate under different circumstances. We offer concrete suggestions for how to make appropriate clinical decisions across a range of circumstances in which a single seizure is observed on routine EEG.

METHODS

Mathematical model for seizure event times and rates

We assume seizures occur at random times, following a Poisson process(Bishop 2006) with an unknown rate, r, between 0 and some maximal physiologically plausible rate R_{max}. Thus the probability that a given number of seizures k occur in a time interval of length T is

$$P(k|r,T)=(rT)^k e^{-rT}/k!$$

In the following sections, we assume that we have observed 1 seizure ($k = 1$) during a "routine" outpatient EEG recording lasting $T = 30$ minutes.

We explore the effects of different possible pre-EEG estimates of seizure rate by examining how different assumptions regarding pretest probability estimates of seizure frequency, denoted $P(r)$, affect post-test estimates of seizure rate, denoted $P(r|k,T)$. The relationship between the pretest probability, $P(r)$, posttest probability, $P(r|k,T)$, and likelihood function, $P(k|r,T)$ is specified by Bayes' rule,

$$P(r|k,T)=P(r)P(k|r,T)/Z$$

where Z is the normalization constant of the distribution. We model prior probability distributions over seizure rates using Gamma distributions(Bishop 2006; El-Sayyad and Freeman 1973),

$$P(r)=\beta^\alpha r^{\alpha-1}e^{-\beta r}/\Gamma(\alpha)$$

where Γ is the normalization constant. The mode and standard deviation of the Gamma distribution are given by: mode = $(\alpha-1)/\beta$, and standard deviation = $\sqrt{\alpha}/\beta$. The Gamma distribution is flexible enough to model a wide range of clinically-relevant scenarios, and has the mathematical advantage that, when paired with the Poisson model for seizure event times, the post-test probability distribution $P(r|k,t)$ is also a Gamma distribution(Bishop 2006). Updated values for the post-test distribution are given by simple formulae(Bishop 2006; El-Sayyad and Freeman 1973), namely: α is updated to $\alpha + k$, and β is updated to $\beta + T$. Specific values of α and β used in the figures, together with the corresponding mean and variance values from which they are computed, are as follows, with values given in the format (mean,standard deviation, α, β): for Figure 2A, (4,1, 17.9, 4.2) and (4.2, 1.0, 18.9, 4.2); and for Figure 2B (4, 8, 1.6, 0.2) and (10.2, 10.1, 2.6, 0.2). As explained below, Figure 1A was constructed assuming a uniform distribution over the range from 0 to 4,3200 seizures / month, coupled with the following values for the likelihood function: $k = 1$, and $T = 30$ minutes.

RESULTS

Scenario #1: Unbiased ignorance

First consider the situation where the EEG reader knows no clinical details of the case, and wishes to "let the data speak for itself". More explicitly, assume that there is no *a priori* reason to favor any particular value in the large range of possible seizure rates (0 through R_{max}), hence prior to seeing the EEG the prior probability distribution is uniform across this range. In this case, Bayes' rule dictates that the shape of the post-test probability distribution is governed entirely by the shape of the likelihood function. We specify the range of "plausible" seizure rates as those for which there is a 95% probability that the true seizure rate lies between them.

This distribution is shown in Figure 1, along with the boundary of the range containing 95% of the probability. We assume a theoretical maximum possible rate of 1 seizure/minute (~43,200/month). The 95%-probability interval in this case of an uninformed prior probability distribution is broad, approximately 120 – 13,778 seizures per month (4–459 seizures / day). As intuition would suggest, the most probable seizure rate (the mode of the distribution) is simply 1 seizure per half hour, or 1,440 seizures per month. Lower rates are of course possible, but less probable. For example, the probability that the true rate is <4 per month would be judged <0.001%.

Scenario #2: Well-known seizure frequency

Next, suppose instead that we have reliable clinical and/or prior EEG information suggesting that the most likely baseline seizure frequency is 4 per month, and that we are fairly confident of this estimate; e.g. suppose the error is not much more than +/− 1 seizure per month. For example, this level of confidence could be afforded by a meticulously seizure diary kept by a reliable reporter, such as a patient whose seizures do not impair awareness, or an observant family member who spends sufficient time with the patient. The resulting probability distribution, depicted by the dashed curve in Figure 2A, serves as the pre-test probability distribution regarding seizure frequency.

We now update this distribution, incorporating the observation of a single seizure in a 30 minute recording, again using Bayes' rule. The solid curve in Figure 2A shows that in this case the post-test probability distribution is not much changed from the pre-test distribution: The mode increases by a clinically insignificant amount, from 4 to 4.24 seizures/month,. Because the seizure frequency is fairly certain prior to testing, the single observed seizure has little impact on our clinical impression of seizure frequency. Bayes' rule helps us reconcile that while the observation is surprising (the probability that a single seizure occurs during any given 30 minute EEG recording is only 0.27–0.29%), this observation is counterbalanced by our strong prior knowledge that the seizure frequency was low.

Scenario #3: Poorly-known seizure frequency

Suppose now that information about the patient's seizure rate is less certain, as when the history is unreliable, we have only known the patient for a brief time, or we suspect the patient is often unaware of seizures(Burneo 2008; Hoppe et al. 2007; Blum et al. 1996; Kerling et al. 2006). Here, we expect a single observed seizure to carry relatively more weight, because the prior knowledge is weak. Suppose our pre-test probability distribution has a mode of 4 seizures / month, and a large standard deviation of, say, 8 seizures / month. This level of (un)certainty might be arrived at, for example, in managing a patient who is not able to reliably report an accurate seizure frequency, but for whom we have arrived at a baseline estimate based on short-term monitoring data collected via ambulatory EEG or admission to an epilepsy monitoring unit. In such a case, a single seizure in a 30 min test

causes us to increase our estimate of the most likely seizure frequency by ~6, from 4 up to 10.1 seizures / month (Figure 2B).

Adjusting frequency estimates after a seizure on routine EEG: A clinical guide

As suggested by our three examples, the degree of uncertainty in our pre-test estimates of seizure frequency determines the amount by which our estimates are influenced by observing a seizure during a brief EEG recording; uncertain pre-test estimates should always be more heavily influenced by new data. Figure 2C illustrates this general relationship, and may serve as a clinical guide to adjusting one's pre-test probability estimate of seizure frequency after seeing a seizure on a brief EEG recording. For example, a pre-test estimate of 2 ± 2 seizures / month (mode ± standard deviation) is adjusted upward to 3.2 ± 2.3, an increase of 1.2; whereas a less-certain pre-test estimate of 2 ± 4 is affected more, increasing to 5.1 ± 5.1. Similarly, a pre-test rate of 20 ± 20 is increased by 12.1, to 32 ± 23.3, whereas a less certain *a priori* estimate of 20 ± 40 increases by more than double, to 50 ± 49.7.

"Don't be fooled by randomness": The fallacy of division

Finally, consider the following line of reasoning, the "don't be fooled by randomness"(Taleb 2001) argument: In any large number of routine EEGs we are bound to see a few seizures "by chance", even if all patients' underlying seizure rates are low. For example, if we perform a routine (30 minute) EEG on 5,000 patients who all have seizures at average rates of 1 per month, then we should expect to see 3–4 seizures among this set of recordings. Thus, a single seizure during a routine EEG recording need not imply a high underlying seizure rate.

While this observation correctly points out that a single seizure does not *necessarily* imply a high underlying seizure rate, it is problematic to conclude that the underlying rate is *probably* actually low in any given individual. While inference is straightforward when specific information about an individual's seizure rate is known (Scenarios 2 above), the most appropriate choice of prior probability distribution when no specific clinical information exists is not obvious. When the population referral base is broad and inhomogeneous, as is often the case, the true underlying distribution of rates for the population may have multiple peaks or "fat tails" (non-negligible probability at high rates); such "fat tailed" distributions are notoriously difficult to accurately measure empirically, and are susceptible to severe underestimation, especially with respect to low probability events(Taleb 2010, Taleb 2001). Thus the "don't be fooled by randomness" approach leaves one vulnerable to discounting outliers who actually do have very high seizure rates due to model inaccuracies.

By contrast, the "unbiased" approach of modeling all possible seizure rates as *a priori* equally probable leads to the conclusion that the underlying seizure rate is likely to be high, though the range of probable rates will be quite broad (see Scenario 1 above). The primary consequence of this approach is to lead one to seek further clinical information, which can either confirm a high seizure rate and lead to appropriate medical intervention, or provide reassurance that the observed seizure was most likely a rare event on a background of infrequent seizures.

The above discussion suggests that the impulse to discount the significance of a seizure during routine EEG based on "don't be fooled by randomness" reasoning is generally a variation on the "fallacy of division": making inferences about individuals by assuming that individual members of a group have the typical characteristics of the group(Robinson 2009).

DISCUSSION

We have seen that a single seizure during a brief routine EEG recording can have very different implications depending on the prior clinical information available. In cases where nothing is known beyond the EEG, a conservative assumption that all physiologically plausible seizure rates are equally likely necessarily leads us to suspect that the underlying seizure rate is quite high. Alternatively, when prior clinical data constrains the true underlying seizure frequency within a narrow range, a single seizure is properly interpreted as a rare event, and estimates of the underlying seizure rate remain heavily determined by the prior clinical knowledge. In intermediate cases, with less certain clinical data, the post-test probability distribution is determined by substantial contributions from both prior information and the EEG data. Finally, we have argued the need for caution in discounting a single seizure in an individual on the grounds that occasional seizures are expected by chance alone in a large set of EEGs, because this reasoning falls prey to fallacy of division.

These considerations lead us to offer the following concrete, situation-dependent recommendations, summarizing appropriate reactions and reasoning in response to a single seizure during a routine recording. The recommendations apply to patients of all ages, and in the final analysis may seem commonsensical. However, these issues have arisen repeatedly over many years in discussions among the authors and colleagues, therefore we believe it is useful to state these recommendations explicitly.

Scenario # 1. No prior knowledge

A single seizure implies a high underlying rate of seizures (hundreds to thousands per month), until proven otherwise. More information should be sought until the true seizure frequency can be estimated with sufficient confidence to allow appropriate clinical action, e.g. by speaking with the patient, family members, caretakers, or referring physician, reviewing prior EEG data, or obtaining a follow-up or longer EEG recording.

Scenario # 2.Tightly-known seizure frequency

If prior information supports a fairly low seizure probability with reasonable certainty, then the single event need not lead to any substantive revisions of clinical impression nor alter previous management decisions. Nevertheless, the sensitive dependence of this conclusion on the prior probability distribution serves as a cautionary reminder to be careful how confident we are about our prior estimates, especially in cases where seizures are likely to be underreported (e.g. seizures which the patient is unaware of are likely to be underreported if they produce subtle manifestations or if the patient lives alone).

Scenario # 3. Loosely-known seizure frequency

This case is intermediate between #1 and #2, and simply requires an intermediate level of intensity of further investigation.

Acknowledgments

This work stems from a series of discussions between the authors and other members of our clinical group including Ioannis Karakis, Ronan Kilbride, Jay Pathmanathan and others. SSC, AJC, DBH and MBW receive support from NINDS–NS062092

References

Bishop, CM. Pattern recognition and machine learning. New York: Springer; 2006.

Blum DE, et al. Patient awareness of seizures. Neurology. 1996; 47(1):260–264. [PubMed: 8710091]

Burneo JG. The real truth behind seizure count. Epilepsy Currents/American Epilepsy Society. 2008; 8(4):92–93. [PubMed: 18596872]

El-Sayyad GM, Freeman PR. Bayesian sequential estimation of a Poisson process rate. Biometrika. 1973; 60(2):289–296.

Hoppe C, Poepel A, Elger CE. Epilepsy: accuracy of patient seizure counts. Archives of Neurology. 2007; 64(11):1595–1599. [PubMed: 17998441]

Kerling F, et al. When do patients forget their seizures? An electroclinical study. Epilepsy & Behavior: E&B. 2006; 9(2):281–285.

Robinson W. Ecological Correlations and the Behavior of Individuals. International Journal of Epidemiology. 2009; 38(2):337–341. [PubMed: 19179346]

Taleb, N. Fooled by randomness?: the hidden role of chance in the markets and in life. New York: Texere; 2001. .

Taleb, N. The black swan: the impact of the highly improbable. 2nd ed.. New York: Random House Trade Paperbacks; 2010.

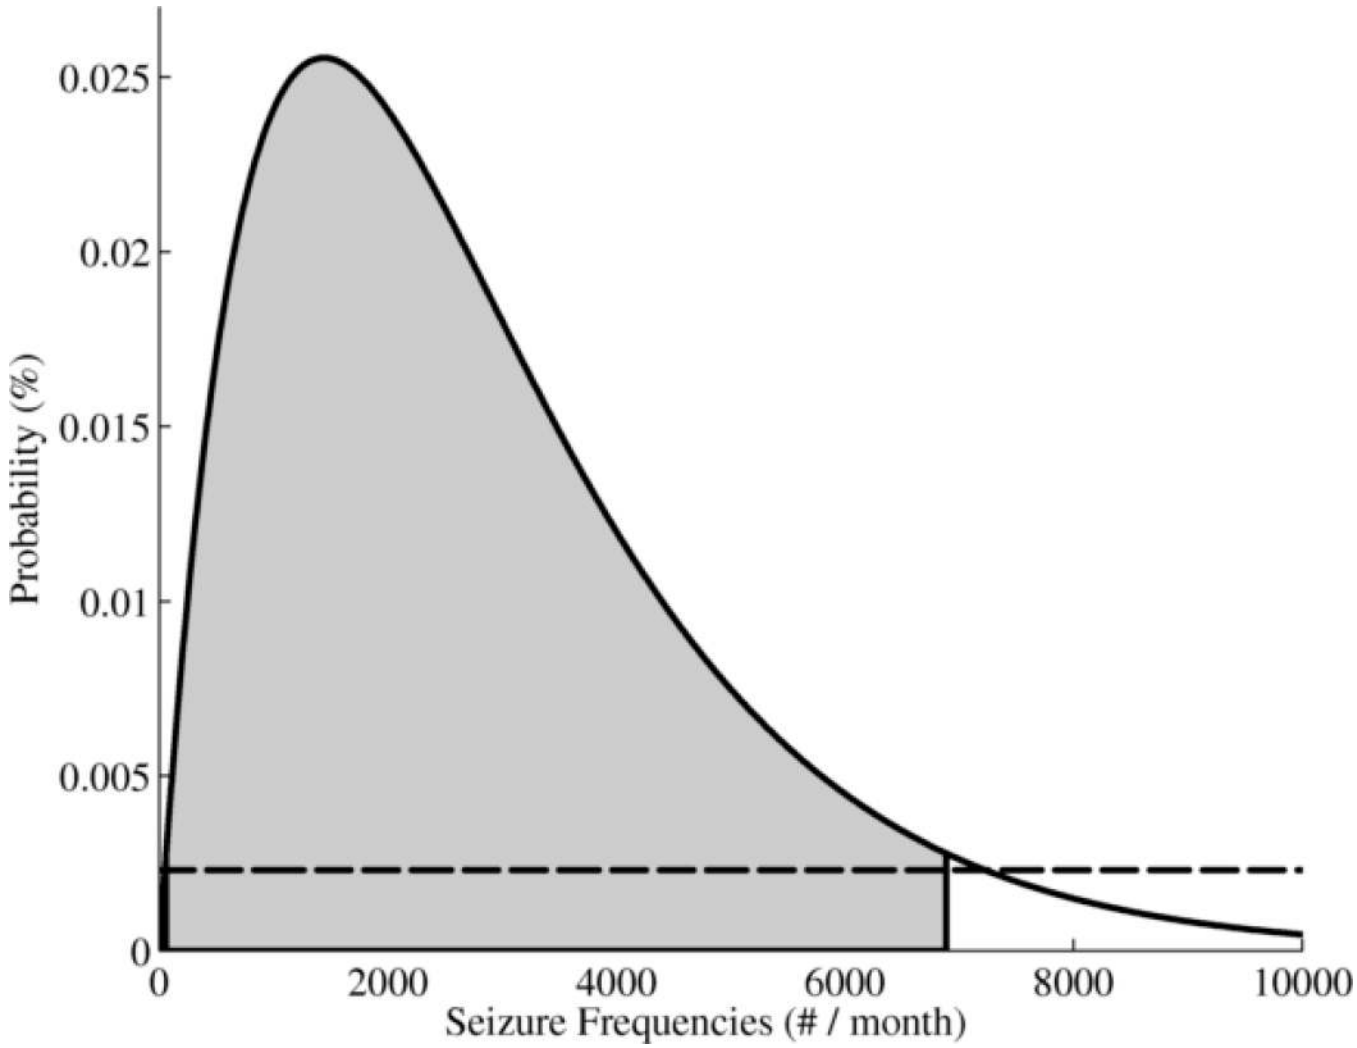

Figure 1.

Pre- and post-test probability distributions for seizure frequency, assuming a uniform pretest probability distribution, with maximum possible rate assumed to be 43,200/month (1 seizure / minute). Dashed line: pre-test probability (uniform distribution); solid line: post-test probability distribution, after observing a single seizure during a routine 30 minute EEG recording. The portion of the distribution containing 95% of the total probability mass is shaded.

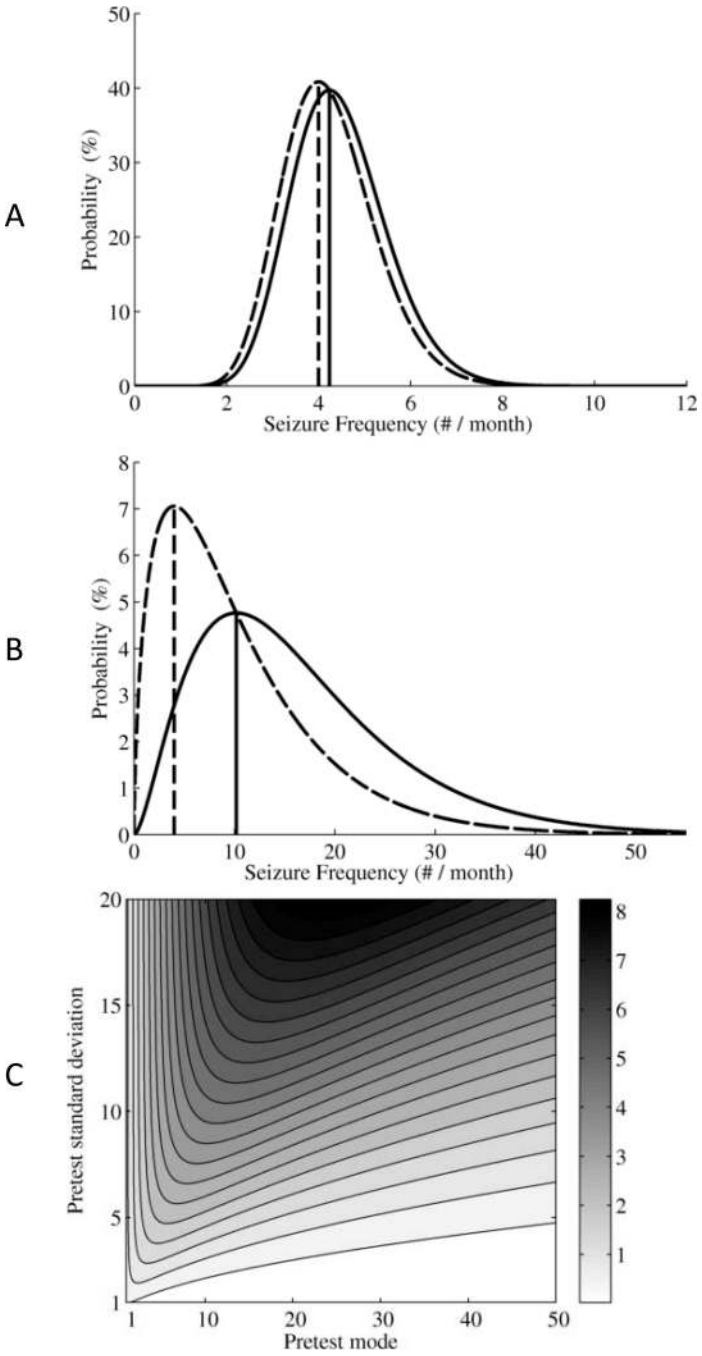

Figure 2.

(A) Pre- and post-test probability distributions for seizure frequency, assuming a tight (narrow) pretest probability distribution. Dashed line: pre-test probability distribution, mode = 4.0, standard deviation (SD) =1; solid line: post-test probability distribution, after observing a single seizure during a routine 30 minute EEG recording, mode 4.24, SD = 1.03. The mode of each distribution is indicated by a vertical line. (B) Same as (A), now assuming loose (broad) pretest probability distribution. Dashed line: pre-test probability distribution, mode = 4, standard deviation (SD) =8; solid line: post-test probability distribution, after observing a single seizure during a routine 30 minute EEG recording, mode 10.2, SD = 10.06. (C) Contour plot illustrating the relationship between uncertainty in the pretest

probability estimate (“pre-test standard deviation”) and the change in the best estimate of the patient’s seizure frequency (mode of the posterior distribution) after observing a single seizure on a 30-minute EEG; the amount of change is encoded in the grayscale. Values are in units of seizures/month.