

OPEN

# A Full-Stack Application for Detecting Seizures and Reducing Data During Continuous Electroencephalogram Monitoring

**BACKGROUND:** Continuous electroencephalogram monitoring is associated with lower mortality in critically ill patients; however, it is underused due to the resource-intensive nature of manually interpreting prolonged streams of continuous electroencephalogram data. Here, we present a novel real-time, machine learning–based alerting and monitoring system for epilepsy and seizures that dramatically reduces the amount of manual electroencephalogram review.

**METHODS:** We developed a custom data reduction algorithm using a random forest and deployed it within an online cloud-based platform, which streams data and communicates interactively with caregivers via a web interface to display algorithm results. We developed real-time, machine learning–based alerting and monitoring system for epilepsy and seizures on continuous electroencephalogram recordings from 77 patients undergoing routine scalp ICU electroencephalogram monitoring and tested it on an additional 20 patients.

**RESULTS AND CONCLUSIONS:** We achieved a mean seizure sensitivity of 84% in cross-validation and 85% in testing, as well as a mean specificity of 83% in cross-validation and 86% in testing, corresponding to a high level of data reduction. This study validates a platform for machine learning–assisted continuous electroencephalogram analysis and represents a meaningful step toward improving utility and decreasing cost of continuous electroencephalogram monitoring. We also make our high-quality annotated dataset of 97 ICU continuous electroencephalogram recordings public for others to validate and improve upon our methods.

**KEY WORDS:** critical care; electroencephalography; epilepsy; machine learning; seizures; software

Continuous electroencephalogram (cEEG) monitoring is an important tool in the ICU for assessing cerebral electrical activity (1) and detecting seizures; however, significant challenges limit its utility as a real-time monitor and hamper more widespread implementation. Continuous electroencephalogram is the test of choice for detecting nonconvulsive seizures including status epilepticus which, in critically ill patients, are both prevalent and associated with mortality (2). Prompt recognition of nonconvulsive seizures is challenging, as data review requires manual interpretation by trained experts, such as physicians or electroencephalogram technologists (3). In addition, demand for cEEG monitoring often fluctuates widely (in our hospital system ranging from one to over 20 cEEG patients monitored at any given time), making staffing challenging and costly (4). Even when such staff are in place, they are limited in the number of patients they are able to monitor due to labor-intensive demands for maintaining the quality of scalp recordings and rapidly interpreting recordings. At many institutions, cEEG is often read

John M. Bernabei, BSE<sup>1,2</sup>  
 Olaoluwa Owoputi, BSE<sup>1,2</sup>  
 Shyon D. Small, BS<sup>1,2</sup>  
 Nathaniel T. Nyema, BS<sup>1,2</sup>  
 Elom Dumenyo, BS<sup>1,2</sup>  
 Joongwon Kim<sup>1,2</sup>  
 Steven N. Baldassano, MD, PhD<sup>1,2</sup>  
 Christopher Painter, MS<sup>1,2</sup>  
 Erin C. Conrad, MD<sup>3</sup>  
 Taneeta M. Ganguly, MD<sup>3</sup>  
 Ramani Balu, MD, PhD<sup>3</sup>  
 Kathryn A. Davis, MD<sup>2,3</sup>  
 Joshua M. Levine, MD<sup>3</sup>  
 Jay Pathmanathan, MD, PhD<sup>3</sup>  
 Brian Litt, MD<sup>1-4</sup>

Copyright © 2021 The Authors. Published by Wolters Kluwer Health, Inc. on behalf of the Society of Critical Care Medicine. This is an open-access article distributed under the terms of the Creative Commons Attribution-Non Commercial-No Derivatives License 4.0 (CCBY-NC-ND), where it is permissible to download and share the work provided it is properly cited. The work cannot be changed in any way or used commercially without permission from the journal.

DOI: 10.1097/CCE.0000000000000476

at scheduled intervals, usually 8–12 hours, which can lead to significant delays in identification of critical events. Automated systems for evaluating cEEG in real time have the potential to recognize actionable events, such as seizures, much more quickly than manual interpretation and at a lower cost per patient, which could expand the use of cEEG in both resource-rich and resource-poor healthcare settings (5).

Visual quantitative electroencephalogram (qEEG) methods have been deployed to reduce the time and cost associated with manual electroencephalogram interpretation (6, 7). The most commonly used qEEG techniques offer near real-time analysis, displaying compressed metrics derived from amplitude and frequency. Changes in these variables can be used to detect seizures and cortical ischemia. However, although qEEG may significantly reduce review time by the clinician, sensitivity for identifying seizures remains low (51–67%) (6). Despite its advantages over inspection of raw waveform data, visual qEEG still requires specialized training and inspection of the entire recording, albeit in a compressed format.

As an alternative to visual qEEG, there are several algorithms for detecting seizures in scalp electroencephalogram. The most widespread is the "Reveal" algorithm from Persyst (Solana Beach, CA), which has a reported clinical sensitivity of 76% with a false-positive rate of 0.11/hr (8), although subsequent studies have shown a significantly higher false-positive rate (9). This level of performance leads to a significant proportion of seizures being missed, as well as a high false alarm burden. Other "nonpatient-specific" algorithms have been reported to perform better than Persyst but were either studied in epilepsy patients with stereotyped seizures (10) or tested on a carefully curated and cleaned dataset. Patient-specific algorithms have the highest level of performance (11) but require clinicians to mark training data for each individual, which, depending on the time to the first event, renders these approaches less practical for deploying rapidly in an ICU. There is a clear clinical need for nonpatient-specific seizure detection algorithms that are highly sensitive and specific, even when applied to artifact-laden heterogeneous data of the sort typically obtained from ICU electroencephalogram. There is also a need for a gold standard, widely available cEEG data set and objective performance criteria for seizure detection algorithms that can be used by experts and the Food and

Drug Administration for benchmarking cEEG analysis tools, similar to what our group has done for benchmarking seizure detection algorithms for intracranial electroencephalogram (12).

In this study, we introduce a novel framework for semiautomated cEEG analysis and data reduction developed using data collected in the ICU. We share our source code and unique dataset, so that others may improve upon our results and methods. Rather than designing an algorithm to replace clinical cEEG review entirely, we use machine learning to perform data reduction with the intent of increasing the speed and decreasing the cost required to accurately evaluate cEEG data for the presence of seizures. Furthermore, our framework includes a data streaming portal that provides simplified yet detailed data to expedite treatment decisions or guide further electroencephalogram review. We aim to establish a path for easy translation of our methods to clinical care, regardless of electroencephalogram hardware, in a way that will permit increased penetration of cEEG monitoring in hospitals worldwide.

## MATERIALS AND METHODS

Our retrospective dataset consists of 97 patients who were treated in ICUs at University of Pennsylvania Health System and who underwent cEEG monitoring between 2017 and 2020. To develop our algorithm, we used cEEG recordings from 27 randomly selected patients with discrete seizures and 50 recordings from consecutive patients without seizures. We also used a completely unseen test set of 10 patients with seizures and 10 seizure-free patients collected after algorithm training. Of the 97 study patients, the mean age was  $57.0 \pm 18.5$  years, and 53 were female. We used the portions of cEEG records that had high-quality markings including onset and offset times of each seizure. Average electroencephalogram recording time was  $3.0 \pm 1.6$  hours per patient. In patients with seizures, the median number was 4 and the median length was 101 seconds. Clinical characteristics of the development and test cohorts are shown in **Table 1**, and metadata on each individual patient including brief descriptors of each patient's seizures can be found in **Supplemental Table S1** (<http://links.lww.com/CCX/A707>). Data were collected through protocols approved by the institutional review board of the University of Pennsylvania

**TABLE 1.**  
**Clinical Characteristics of Continuous Electroencephalogram Patients**

	Training Set	Test Set
Total number of patients	77	20
Number of female patients	42	11
Age, mean $\pm$ SD	57.6 $\pm$ 18.0	54.8 $\pm$ 21.0
Number of patients with seizures	27	10
Total seizures	265	27
Reason for study, <i>n</i>		
Altered mental status	8	5
Witnessed or reported seizure	15	7
Sepsis/toxic/metabolic disorder	11	2
Intracranial hemorrhage	18	3
Neoplasm	9	0
Anoxic brain injury	4	2
Other/unspecified coma	13	1

We used records from 77 patients in the ICU for algorithm cross-validation and training, and a held-out test set of 20 patients who underwent continuous electroencephalogram after algorithm development. The reason for ordering the study was retrospectively collected from the electronic health record as the most direct factor necessitating the study.

(numbers 820595: “An automated platform for ICU EEG monitoring and visualizing results” and 832104: “Center for Neuroengineering and Therapeutics scalp EEG repository”). Electroencephalogram signals were recorded and digitized at 256 Hz using Natus Xltek equipment (Natus Medical, Pleasanton, CA). All electroencephalograms were acquired using electrodes placed in the international 10–20 configuration. Electroencephalogram recordings were annotated by board-certified clinical neurophysiologists to include times of onset and offset for all seizures. We stored electroencephalogram recordings on <http://ieeg.org> (13), a cloud platform for storing and sharing electrophysiologic data.

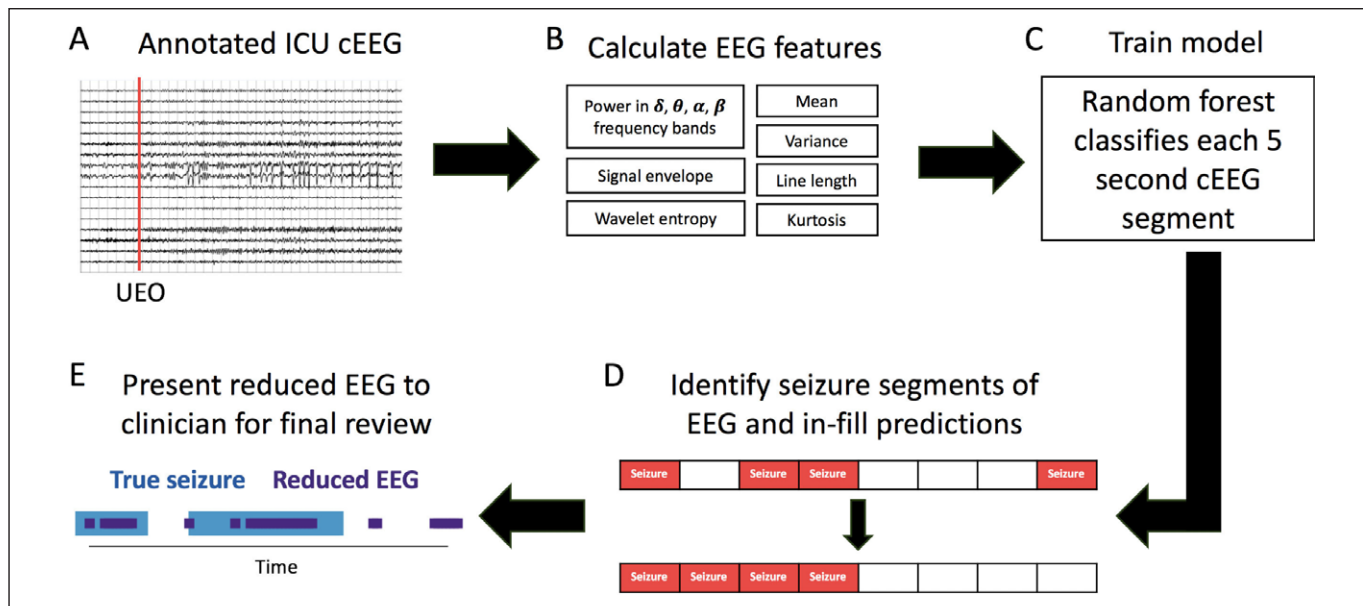
### Feature Extraction

We digitally filtered raw electroencephalogram signals using a fifth-order Bessel bandpass filter with lower and upper cutoff frequencies of 1 and 20 Hz, respectively and calculated features within a nonoverlapping sliding 5-second window. We calculated the following

features for each of the 18 channels in each 5-second window: 1) power in the delta (1–4 Hz), 2) theta (4–8 Hz), 3) alpha (8–12 Hz), and 4) beta (12–25 Hz) frequency bands, 5) signal line length (14), which quantifies the distance between successive points and has been shown to be an effective feature in seizure detection, 6) wavelet entropy (15), which measures the signal complexity in the time and frequency domains and has proven to be an effective electroencephalogram feature (16), statistical features of the signal including the 7) mean, 8) variance, 9) kurtosis, and 10) the mean value of the upper signal envelope of the electroencephalogram waveform. For each feature within each 5-second window, we used both its median value and variance across all electroencephalogram channels yielding 20 total features. Additional detail on feature calculation is found in the supplemental materials (**Methods S1**, <http://links.lww.com/CCX/A707>). Within each window, we used an automated artifact rejection algorithm to remove channels containing missing values or supraphysiologic amplitudes that were clearly due to noise and also excluded any 5-second window with at least three channels containing missing values or shared artifacts that would introduce error into algorithm. At the beginning of feature calculation during model training, the artifact rejection algorithm begins with conservative threshold values of each feature and iteratively rejects segments that surpass those feature levels and checks whether any of the rejected segments were clinically labeled seizures. If so, the threshold of each feature for artifact rejection is raised 50% and the process is repeated, yielding criteria which will not incorrectly reject seizure as artifact in any training patients.

### Machine Learning Approach

We implemented a machine learning framework that identifies electroencephalogram segments of high seizure likelihood in unseen patients. Our algorithmic approach is summarized in **Figure 1**. We used seizures in which the unequivocal seizure onset and offset on electroencephalogram are marked by board-certified electroencephalogram readers for clinical purposes using the method of Litt et al (17) (Fig. 1A). We selected a random forest approach as it is fast, robust to overfitting, and can handle noisy data including correlated features that



**Figure 1.** Seizure detection and data reduction methods. **A**, We use clinically annotated ICU continuous electroencephalogram (cEEG). **B**, We calculate the listed electroencephalogram (EEG) features for each channel subtracted from a common average reference before taking the mean and variance of each feature across channels. **C**, We train a random forest model to classify each 5 s cEEG segment as likely or unlikely to contain seizure and test in unseen patients. **D**, We smooth predictions to improve interpretability for future clinical review in **(E)**. UEO = unequivocal electrographic onset.

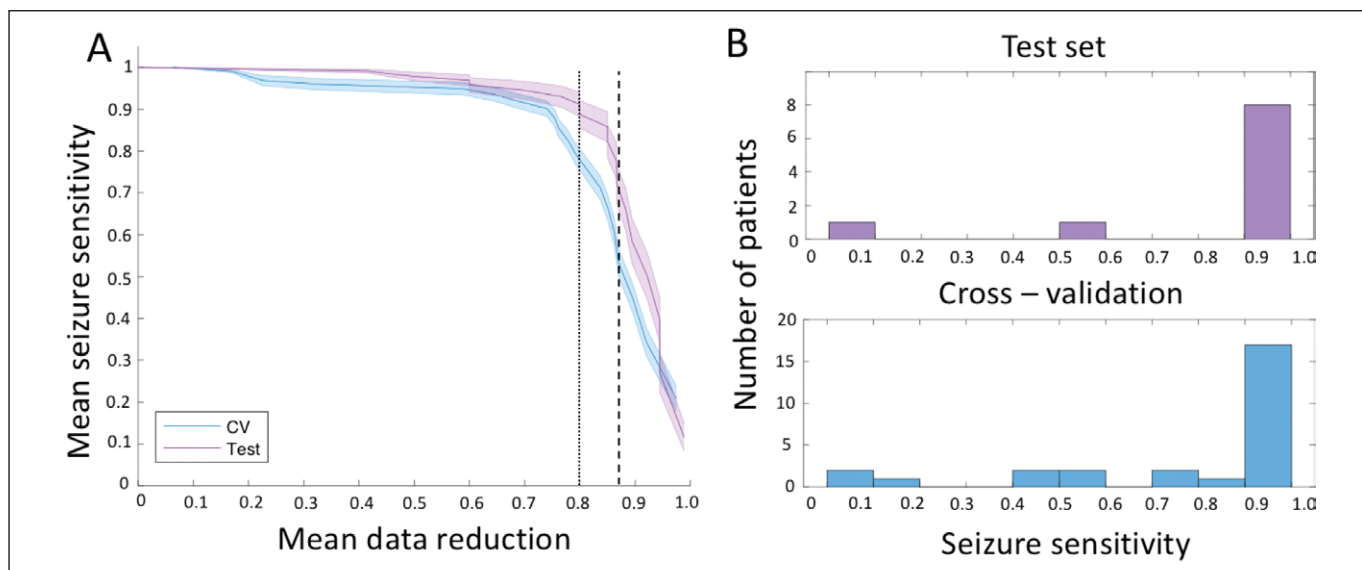
is often found in biomedical applications (18). We trained the model (Fig. 1C) using 400 trees to predict whether each 5-second electroencephalogram segment in each patient contains a seizure or not based on the 20 features discussed in the prior section (Fig. 1B). As seizures are typically uncommon events, even in recordings that contain them, they make up only a small proportion of our overall dataset. Thus, we train the classifier to be penalized 500 times as heavily for false negatives (e.g., missing a seizure) as for false positives. In algorithm development, we used five-fold cross-validation in which one of five patients were held out of training at a time to be used for validation. To quantify final performance, we trained the algorithm on all 77 development patients and tested on a separate cohort of 20 patients. After the model generates its prediction of seizure or nonseizure for each 5-second window, the predictions are then briefly postprocessed to improve readability and clinical workflow. For any single 5-second windows that our system deemed “nonseizure” but is surrounded on either side by windows marked as “seizure,” we change the marking of the center window to “seizure,” so that the reduced electroencephalogram becomes more continuous (Fig. 1D). We subsequently remove “seizure” markings initially made by our system

which are not within 15 seconds of another “seizure” marking as this lowers our chance of false positives. These steps make the reduced electroencephalogram significantly more contiguous and amenable to clinical review.

## Evaluation

To measure the performance of our model, we calculated “seizure sensitivity” and specificity, which indicates the level of “data reduction” our system can achieve. We define “seizure sensitivity” as the proportion of seizures for which the algorithm marks either the entirety or a portion of the event on electroencephalogram. We calculate specificity and thus “data reduction” by determining the proportion of true-negative windows in the patient’s time series that are marked for removal rather than review by the clinician. Our random forest classifier outputs the probability of seizure for each window and adjusting the threshold from its default of 0.5 (e.g., any 5 s window with > 50% seizure likelihood is marked for clinician review) allows us to generate receiver operating characteristic curves to assess the tradeoff between seizure sensitivity and data. Each clinician user could adjust this threshold to tune the seizure detection-data reduction tradeoff to their own preferences.





**Figure 2.** Algorithm performance. **A**, Tradeoff between seizure sensitivity and mean data reduction for both the cross-validation (CV) (blue) and test (purple) sets. Vertical lines: default operating performance defined by a seizure classification threshold of 50% (blue dotted line at 83% specificity for CV, purple dashed line at 89% specificity for the test set). **B**, Histogram of seizure sensitivities at the default operating points for both the CV set (bottom, blue), and the test set (top, purple).

## Full-Stack Application

We construct an integrated application to manage the inputs, outputs, data storage of our novel machine learning algorithm, and its interaction with users. In our application, an open-source task management platform called “Celery” allocates separate, asynchronous processes which harvest and process data from electroencephalogram streams, calculate seizure likelihood using the algorithm, and ultimately store predictions in a MongoDB database (**Supplementary Fig. S1A**, <http://links.lww.com/CCX/A707>). In this study, we used <http://ieeg.org> on Amazon’s elastic computing cloud for electroencephalogram storage and its toolboxes for data streaming into custom MATLAB software (MathWorks, Natick, MA). Implementation could also be performed on local machines behind institutional firewalls or on HIPAA compliant cloud facilities, as optimal for individual institutions. We also present an interactive, web platform using the Python Flask library to display reduced electroencephalogram and allow clinicians to interact with and understand the outputs of our system (**Supplementary Fig. S1, B and C**, <http://links.lww.com/CCX/A707>).

The main page of the dashboard shows an overview of all patients who are currently undergoing cEEG monitoring in the ICU with the real-time, machine learning–based alerting and monitoring system

for epilepsy and seizures (RAMSES) system. Each patient is listed with information related to outputs from the classifier. We show the number of seizures detected over the course of a patient’s recording, the percentage of the recording consisting of concerning epochs, the time in minutes since the last seizure, and a visual representation of the most concerning prediction over the length of the recording. This visual representation appears as a dot color coded according to the respective prediction as follows: red is a likely seizure, yellow is a potential seizure, and blue is non seizure. Clinicians may modify the dashboard layout and display statistics according to their preferences. Specifically, they may order patients according to those with most recent seizures, highest density of concerning epochs, or the room number. Furthermore, they may adjust the time period over which the statistics are calculated to provide different quantifications of clinical status.

Within the same application, clinicians may also select any given patient and further inspect the algorithm’s outputs over the duration of the recording. In this patient-specific view, predictions are represented as a timeline with different epochs color coded in accordance with the color associations of the dots on the main page. Clinicians can further inspect the raw electroencephalogram associated with each prediction by double-clicking the prediction on timeline.

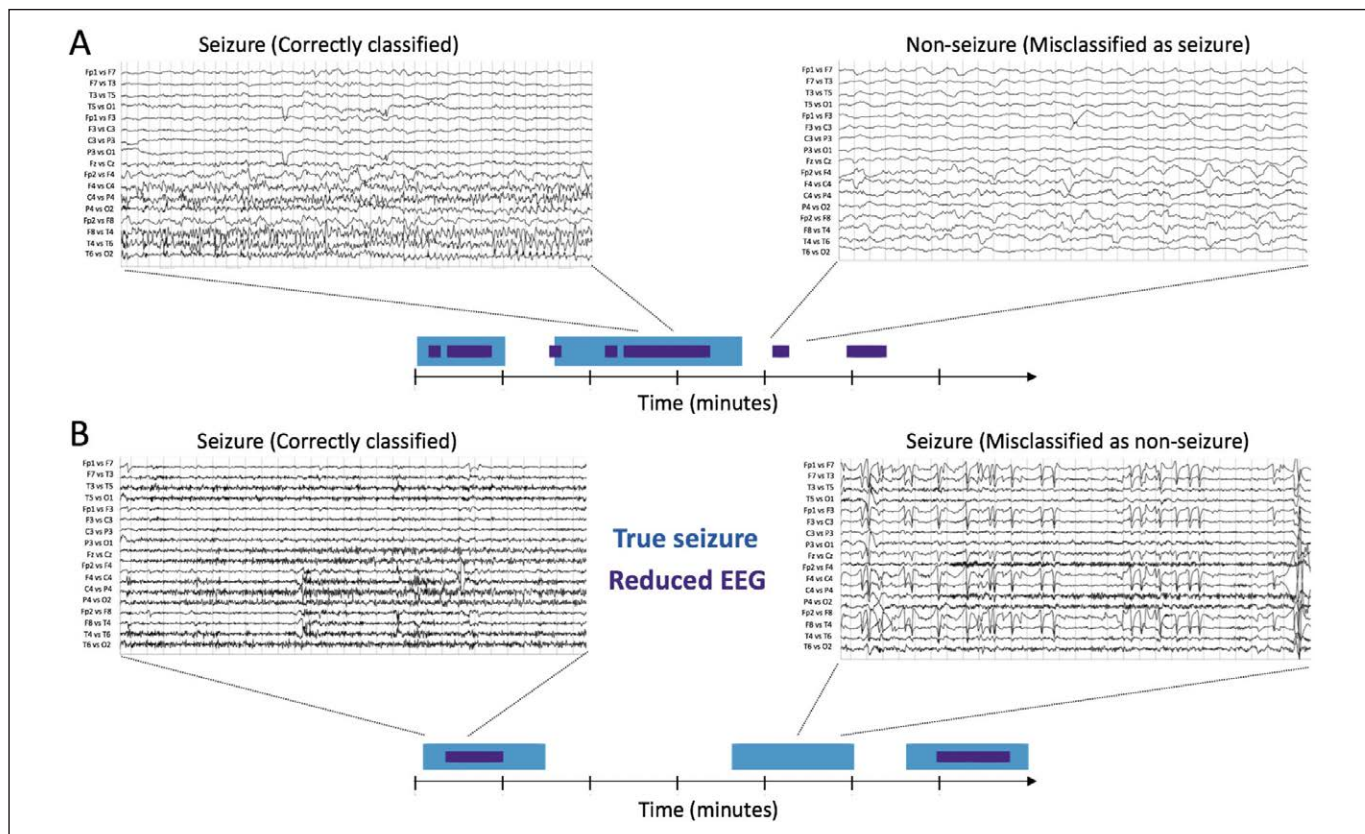
## Data and Code Sharing

All records and annotations are freely available on <http://ieeg.org> associated with the patient IDs listed in Supplementary Table S1 (<http://links.lww.com/CCX/A707>). The code for the seizure detection and data reduction algorithms is available at [GitHub.com /jbernabei/ICU\\_EEG](https://github.com/jbernabei/ICU_EEG), whereas the code for streaming and web-interfacing is available at [GitHub.com/nathanielnyema/RAMSES](https://github.com/nathanielnyema/RAMSES). We aim for our methods to be translatable and for other groups to validate and improve our algorithms or their own using the resources we provide.

## RESULTS

We retrospectively acquired data from 97 critical care patient cEEG recordings including individuals with and without seizures. Our artifact rejection process yielded a mean of 44 rejected 5-second intervals; however, the range across patients was high with 61 of 97 patients having no rejected windows and five of 97

patients having over 200 rejected windows, mirroring the wide variability in clinical care and sources of potential artifacts in the ICU. To assess seizure detection performance and its trade off with the amount of data reduction, we calculated five-fold cross-validation seizure sensitivity across our development cohort of 77 patients. **Figure 2A** shows the ROC curve representing the performance of our system for both the cross-validation (blue) and test (purple) sets across a range of classification thresholds from 5% to 100% seizure likelihood. At our default setting of 0.5 representing the truest performance of our classifier (e.g., 5 s windows with > 50% likelihood of containing seizure activity are classified as such), we found a mean seizure sensitivity of 84% in the cross-validation and 85% in the test sets, respectively. The exact distributions are found in **Figure 2B**, and the number of seizures missed in each patient are noted in Supplementary Table S1 (<http://links.lww.com/CCX/A707>). In the patients within the cross-validation and test sets which contain seizures, specificity and thus data reduction was 80% (range, 6–100%;



**Figure 3.** Representative results of data reduction algorithm. In both panels, the distribution of true seizures over an 8-min period are shown in *blue*, and the reduced electroencephalogram (EEG) is shown in *purple*. All EEG is displayed in anterior-posterior bipolar montage and is of 35 s in length. **A**, Continuous EEG (cEEG) clip of a true-positive (*left*) and false-positive (*right*) seizure segments. **B**, cEEG clip of true-positive (*left*) and false-negative (*right*) seizure segments.

interquartile interval 77–94%) and 87% (range 46–100%; interquartile interval 83–99%), respectively. For seizure-free patients, comprising 50 of 77 cross-validation subjects and 10 of 20 test subjects, the average specificity and thus data reduction was 84% (range 14–100%; interquartile interval 79–97%) and 85% (range 41–99%; interquartile interval 79–97%) for the cross-validation and test sets, respectively. Of the 27 cross-validation patients and 10 test set patients who had seizures during the recording period, the data reduction algorithm only missed all seizures in one patient in each set. In the cross-validation set, the patient had a single 10-second event lacking high-frequency activity over a low-voltage background (**Supplementary Fig. S2A**, <http://links.lww.com/CCX/A707>), whereas in the test set, the patient had seizure activity evident only in a small number of channels superimposed over higher amplitude spikes present at baseline throughout much of the record (**Supplementary Fig. S2B**, <http://links.lww.com/CCX/A707>). In all other patients, few seizures were missed. To quantify our performance at an additional operating point, we adjusted the classification threshold to favor greater than 90% mean sensitivity across patients. This step resulted in 75% data reduction in the cross-validation set but only 55% data reduction in the test set.

To provide a better understanding of the strengths and limitations of a clinical implementation of our system, we visualize examples of system outputs in **Figure 3**. **Figure 3A** shows true-positive and false-positive electroencephalogram examples for a patient in which the algorithm correctly identified 13 of 18 seizures while achieving a 99.2% data reduction. On the left, we show a clip of correctly classified seizure activity localized to the right hemisphere, whereas on the right, we show a nonseizure segment which the algorithm erroneously classified as possibly containing a seizure. The asymmetry of right and left hemisphere activity in the false-positive example could have skewed our algorithm to predict seizure during this time segment. **Figure 3B** shows a patient in which our algorithm correctly identified 36 of 40 seizures with a 91.7% data reduction. On the left, a correctly identified seizure is displayed, whereas on the right, there is an example of a missed seizure where strong discharges in the frontal electrodes may have masked the high-frequency activity in feature space. This algorithm achieved a seizure detection sensitivity of 90% for this patient.

## DISCUSSION

In this study, we present an important step in developing and implementing automated cEEG analysis systems to manage the increasing demand for expensive ICU electroencephalogram monitoring. Our main objective was to use machine learning to both reliably detect seizures and to dramatically reduce the amount of electroencephalogram that must be physically reviewed by physicians and trained technologists. Additionally, we aimed to provide an open-source framework to allow data handling, storage, and display—which could be applied to other uses of electroencephalogram monitoring. We found that this approach provides a mean seizure sensitivity of 84% in cross-validation and 85% in testing, as well as a mean specificity of 83% in cross-validation and 86% in testing. In both cross-validation and testing, the majority of patients had all seizures detected (19/27 for cross-validation, 8/10 for test). We share all of our data and code with the intention that our methods are improved upon, so that machine-learning assisted data reduction can be used clinically to expand the use and decrease the cost of continuous electroencephalogram in the ICU setting.

Our algorithm is primarily distinct from others by its approach as data reduction rather than an expert labeler of seizure onset and offset or of patients that contain seizure versus those which do not. Rather, we draw attention to concerning segments in each record and attempt to minimize the number of seizures that would be missed by implementing such a system. As such, even the majority of patients without seizures have some of their record highlighted and marked for clinician review. These segments could correspond to cerebral activity on the ictal-interictal continuum, or periodic discharges which are abnormal and could be important to clinicians, but do not meet criteria for seizure.

A key question in studies of this nature is what analysis performance metrics are adequate for clinical deployment. Seizure sensitivity is the typical gold standard, although seizure labeling varies significantly between experts (19). Furthermore, the open nature of our processing algorithm allows our algorithm to train toward any given experts reporting style. There is also some indication that to adequately manage patients, it may not be necessary to capture and identify every seizure on electroencephalogram, as many of these events may not have



clinical significance, and trends in number of events, combined with clinical metadata may be adequate for excellent patient management even if a small percentage of subtle electrographic seizures are missed by the algorithm. Indeed, in our results, there are a small number of patients with seizures (1/10 in our test set) for which the algorithm misses all events in that patient. It is likely, as we highlight in Figure 3 and Supplementary Figure S2 (<http://links.lww.com/CCX/A707>) and **Supplementary Figure S3** (<http://links.lww.com/CCX/A707>) that the reasons our algorithm misses seizures could be similar to those that cause humans to miss seizures: subtle features when compared with background, confusion by artifact or noise, and ambiguity of whether the data meet clinical criteria for a seizure. Although there are no studies of the sensitivity of real-time human reviews of multipatient cEEG data, our experience suggests it is likely similar to that of our algorithm.

Although very encouraging, our study has limitations. One limitation is the relatively small sample size of patient data for the purposes of cross-patient seizure detection. In our cohort, 37 of the 97 patients had seizures, which likely did not contain sufficient variety to fully represent the myriad seizure types and locations encountered in an ICU population. This limitation restricts the types of features and models that we can use to those which perform well on small amounts of data. Another potential limitation is the high prevalence of patients with seizures in our dataset compared the general ICU population. However, the similar data reduction performance between patients with and without seizures (87% vs 85%, respectively) implies that there should not be significant variability based upon the proportion patients with seizures. Namely, our study suggests that clinicians using our system would have to spend approximately 85% less time reading each cEEG record to either identify seizures or to be confident that it is unlikely to contain them. Furthermore, the algorithm may not retain the entirety of each seizure in the “reduced” dataset. As each window is 5 seconds in length, portions of the seizure such as the beginning in which a window may overlap both ictal and preictal data could erroneously be marked as normal, as could the initial segments of long seizures with gradual onset in which quantitative features have not yet significantly changed. However, we feel that these seizure clips would be easy for clinicians to notice in the reduced cEEG, and their prompt recognition could outweigh the inconvenience

of failing to capture some events in their entirety. Finally, we do not have training labels for other clinically important phenomena such as sleep stages or different types of nonictal discharges that may influence classification. Detection of such interictal abnormalities is poor with all commercially available software. Despite these shortcomings, we believe that RAMSES is both novel and an important step in moving toward automated monitoring systems that can be rapidly implemented, reduce cost, and increase the efficiency of clinicians and technologists in busy cEEG monitoring settings.

## CONCLUSIONS

The RAMSES system represents a significant starting point for future work in data-driven ICU electroencephalogram analysis. We have previously reported an integrated data management and caretaker notification platform for multimodal ICU data (20). Our vision involves a unifying data platform that is capable of incorporating any number of analytic engines, harnessing the power of cloud computing, and providing real-time clinical updates. Indeed, implantable devices for seizure detection and stimulation may operate under such a paradigm in the near future (21). At present, the modular structure of RAMSES ensures that these future iterations can be incorporated without disruption of the system. It is reasonable to envision this data set dramatically expanding, as we set up multi-institutional collaborations to enhance our performance, data acquisition, and system testing. Our ultimate goal is for many more critically ill patients to benefit from continuous electroencephalogram monitoring during their hospitalizations, while reducing the costs and improving the effectiveness of these systems.

## ACKNOWLEDGMENT

We would like to thank Jacqueline A. Boccanfuso and Amanda Samuel for their assistance in data acquisition.

- 1 Department of Bioengineering, University of Pennsylvania, Philadelphia, PA.
- 2 Center for Neuroengineering and Therapeutics, University of Pennsylvania, Philadelphia, PA.
- 3 Department of Neurology, Hospital of the University of Pennsylvania, Philadelphia, PA.
- 4 Department of Neurosurgery, Hospital of the University of Pennsylvania, Philadelphia, PA.



Supplemental digital content is available for this article. Direct URL citations appear in the printed text and are provided in the HTML and PDF versions of this article on the journal's website (<http://journals.lww.com/ccejjournal>).

Dr. Bernabei acknowledges support from the National Institutes of Health (NIH) 6T32NS091006. Dr. Davis acknowledges support from an NIH K23-NS073801 and the Thornton Foundation. Dr. Litt acknowledges support from the Penn Center for Healthcare Innovation, the Mirowski Family Foundation, and an NIH DP1: NS122038 (principal investigator: Dr. Litt).

For information regarding this article, E-mail: [John.Bernabei@pennmedicine.upenn.edu](mailto:John.Bernabei@pennmedicine.upenn.edu)

## REFERENCES

- Scheuer ML: Continuous EEG monitoring in the intensive care unit. *Epilepsia* 2002; 43(Suppl 3):114–127
- Kubota Y, Nakamoto H, Egawa S, et al: Continuous EEG monitoring in ICU. *J Intensive Care* 2018; 6:39
- Herman ST, Abend NS, Bleck TP, et al: Consensus statement on continuous EEG in critically ill adults and children, Part I. *J Clin Neurophysiol* 2015; 32:87–95.
- Young GB: Continuous EEG monitoring in the ICU: Challenges and opportunities. *Can J Neurol Sci* 2009; 36(Suppl 2):S89–S91
- Baldassano SN, Hill CE, Shankar A, et al: Big data in status epilepticus. *Epilepsy Behav* 2019; 101:106457
- Haider HA, Esteller R, Hahn CD, et al: Critical Care EEG Monitoring Research Consortium: Sensitivity of quantitative EEG for seizure identification in the intensive care unit. *Neurology* 2016; 87:935–944
- Stewart CP, Otsubo H, Ochi A, et al: Seizure identification in the ICU using quantitative EEG displays. *Neurology* 2010; 75:1501–1508
- Wilson SB, Scheuer ML, Emerson RG, et al: Seizure detection: Evaluation of the Reveal algorithm. *Clin Neurophysiol* 2004; 115:2280–2291
- Kelly KM, Shiau DS, Kern RT, et al: Assessment of a scalp EEG-based automated seizure detection system. *Clin Neurophysiol* 2010; 121:1832–1843
- Orosco L, Correa AG, Diez P, et al: Patient non-specific algorithm for seizures detection in scalp EEG. *Comput Biol Med* 2016; 71:128–134
- Park Y, Luo L, Parhi KK, et al: Seizure prediction with spectral power of EEG using cost-sensitive support vector machines. *Epilepsia* 2011; 52:1761–1770
- Baldassano SN, Brinkmann BH, Ung H, et al: Crowdsourcing seizure detection: Algorithm development and validation on human implanted device recordings. *Brain* 2017; 140:1680–1691
- Kini LG, Davis KA, Wagenaar JB: Data integration: Combined imaging and electrophysiology data in the cloud. *Neuroimage* 2016; 124:1175–1181
- Esteller R, Echaz J, Tchong T, et al: Line length: An efficient feature for seizure onset detection. *Conf Proc IEEE Eng Med Biol Soc.* 2001; 2:1707–1710
- Rosso OA, Blanco S, Yordanova J, et al: Wavelet entropy: A new tool for analysis of short duration brain electrical signals. *J Neurosci Methods* 2001; 105:65–75
- Kannathal N, Choo ML, Acharya UR, et al: Entropies for detection of epilepsy in EEG. *Comput Methods Programs Biomed* 2005; 80:187–194
- Litt B, Esteller R, Echaz J, et al: Epileptic seizures may begin hours in advance of clinical onset: A report of five patients. *Neuron* 2001; 30:51–64
- Boulesteix A-L, Janitza S, Kruppa J, König IR: Overview of random forest methodology and practical guidance with emphasis on computational biology and bioinformatics. *WIREs Data Mining Knowl Discov.* 2012; 2:493–507
- Halford JJ, Shiau D, Desrochers JA, et al: Inter-rater agreement on identification of electrographic seizures and periodic discharges in ICU EEG recordings. *Clin Neurophysiol* 2015; 126:1661–1669
- Baldassano S, Gelfand M, Bhalla PK, et al: IRIS: A modular platform for continuous monitoring and caretaker notification in the intensive care unit. *IEEE J Biomed Heal Informatics* 2020; 24:2389–2397
- Baldassano S, Zhao X, Brinkmann B, et al: Cloud computing for seizure detection in implanted neural devices. *J Neural Eng* 2019; 16:026016