My Desktop
Prepare & Submit Proposals
Prepare Proposals (Limited proposal types)
Demo Site: Prepare Proposals
Prepare Proposals in FastLane
Proposal Status
Awards & Reporting
Notifications & Requests
Project Reports
Award Functions
Manage Financials
Program Income Reporting
Grantee Cash Management Section Contacts
Administration
Lookup NSF ID

Preview of Award 1827565 - Annual Project Report

Cover |
Accomplishments |
Products |
Participants/Organizations |
Impacts |
Changes/Problems
| Special Requirements

Cover

Federal Agency and Organization Element to Which Report is Submitted:	4900
Federal Grant or Other Identifying Number Assigned by Agency:	1827565
Project Title:	PFI-TT: Software for Automated Real-time Electroencephalogram Seizure Detection in Intensive Care Units
PD/PI Name:	Iyad Obeid, Principal Investigator Joseph Picone, Co-Principal Investigator
Recipient Organization:	Temple University
Project/Grant Period:	08/01/2018 - 01/31/2021
Reporting Period:	08/01/2019 - 07/31/2020
Submitting Official (if other than PD\PI):	N/A
Submission Date:	N/A
Signature of Submitting Official (signature shall be submitted in accordance with agency specific instructions)	N/A

Accomplishments

*** What are the major goals of the project?**

There are three main goals of this project: **(1) technology hardening:** create a real-time system that is capable of being deployed into clinical environments and support clinical testing, **(2) technology enhancement:** close the gap on performance between our state of the art system and clinically acceptable performance, and **(3) technology evaluation:** evaluate the system on a previously unseen data set to demonstrate that performance translates to a broad range of clinical operating conditions. The expected impact of these three goals is that the research system developed previously under SBIR funding should be transformed into a commercially viable product.

Regarding the first goal, significant challenges include transitioning to a streaming interface with low latency. Research systems tend to be implemented as non-real-time systems that have the advantage of being able to look deep into the past and far into the future to make intelligent decisions. Such systems, including our research prototype, essentially have infinite latency. Hence, our objective is to match our state-of-the-art performance with a system that has a small amount of latency (on the order of seconds) so that it can be clinically viable. The two modules that will need the most work to accomplish this are feature extraction and postprocessing.

Note that not all applications require this. Many Internet-based applications use client-server architectures where the signal is collected in its entirety on a local device before it is processed. As long as back-end processing occurs relatively quickly, users are fine with this type of implementation. However, EEGs can last as long as 72 hours. Neurologists need results immediately after a significant event, such as a seizure, has occurred. Hence, latency is a major issue in this application.

Regarding the second goal, improving sensitivity while maintaining a very low false alarm rate is very important. False alarms are caused primarily by two type of events: artifacts and slowing. A major challenge here will be to implement artifact reduction as a preprocessor. Another challenge will be to use long-term context to disambiguate slowing from a seizure event.

Regarding the third goal, we planned to do two things: (1) develop a blind evaluation set based on TUH EEG data, and (2) integrate an independent evaluation set based on the Duke University Corpus. We will use these corpora for evaluation purposes only, and not optimize the system based on the results on these corpora. We will demonstrate that performance on TUH EEG translates to previously unseen data.

*** What was accomplished under these goals and objectives (you must provide information for at least one of the 4 categories below)?**

Major Activities:

There are three major activities associated with this project: (1) improve the performance of our seizure detection system to meet cinical targets; (2) create a real-time implementation of our best research system that has low latency and is acceptable for clinical use; and (3) continue our attempts to commercialize the technology, building on the results of our STTR Phase I grant.

We began this project with a baseline system that was complex and had significant latency. The system utilized a convolutional neural network (CNN) in its initial stages, then used a long short-term memory network (LSTM) in its latter stages. It also used a heuristic postprocessor. The combination of these things gave us decent performance, but significant latency.

Our first attempt to close the performance gap involved converting this system to a channel-based LSTM system with a heuristic postprocessor. This system improved performance considerably, but still had a significant amount of latency. Most of the latency arose from the postprocessor, which needed to look across long time periods to make intelligent decisions. Postprocessing contributed to a significant portion of the performance and it had to be manually optimized, making somewhat of a support nightmare.

After many attempts to improve this system, we began exploring the use of pretrained networks using an approach known as a residual neural network (ResNet). The use of pretrained networks has grown considerably in the past year as major research organizations like Google have been making models available that are trained on massive amounts of data. We selected an image processing model and decided to transform the sequence to sequence EEG decoding problem into an image processing problem by converting the multichannel signal into an image. This allowed us to leverage available pretrained image processing systems.

This ResNet-based system is relatively simple, processing 5 second frames of a channel downsampled from 250 Hz to 50 Hz. The multiple channels are built into an image so that the decoder sees all channels at once. We believe this closely emulates the way neurologists interpret waveforms. We process data samples directly and no longer need to do a complex feature extraction (which adds latency). There is a 5.12 sec delay for the main part of the system, and a simple heuristic postprocessor follows that adds another 5 secs of delay. The performance of this system is very close to the channel-based LSTM system, but it is a much simpler and less computationally demanding approach.

A DET curve is attached (fig_01.pdf) that compares these systems. The real-time system's performance is reasonably close to the best research system at operating points of interest for critical-care clinical applications.

We also spent a significant amount of time developing a low-latency real-time system. We first focused on our best research system at the time - the channel-based LSTM system. We converted each module to a real-time system, and reformulated the heuristic postprocessor to use less data from the future, thereby lowering the delay. An analysis of the latency is shown in the second attachment (fig_02.pdf). The overall system has 11 seconds of delay, but suffers from a decrease in sensitivity and an increase in the false alarm rate. After several attempts to recover this loss in performance, we decided it was better to switch to the ResNet approach, which drew some of its inspiration from the 2020 Neureka Epilepsy Challenge.

In February of 2020, we were approached by a colleague, Yannick Roy, who runs an organization known as NeuroTechX that claims to be the world's largest collaboration of scientists interested in EEG processing. It claims over 2,300 members. This coupled with the 3,500 susbcribers of our resources gives us a very comprehensive window into the field. NeuroTechX was collaborating with Novela Neurotech, a start up company that put up the prize money ($5,000 for first place, $2,000 for second place), on an open source Kaggle-style challenge. We joined this collaboration and contributed our data and expertise in scoring. Competitors were given TUSZ v1.5.2 training and development data, and were evaluated using a blind evaluation set that we withheld for competitions such as these.

We also developed a scoring paradigm that used our open source scoring software and applied a post-scoring weighting function that emphasized a reduction in false alarms, since these are critical to clinical acceptance. Contests were given the data in early April and given one month to return their results. We handled all the scoring and publication of the results. To collect the prize money, users had to publish their system descriptions.

The winning systems were the systems that judiciously supressed false alarms and optimized their outputs based on the specific weighting function we employed. At the request of the sponsors, we also weighted the number of channels used as part of the metric - the fewer the channels the better. This was a result of the sponsor's interest in low-cost consumer grade technology.

We were also responsible for analyzing the submissions and verifying that they met the requirements of the competition. This gave us a unique opportunity to study best practices.

The top two winning systems performed at levels very similar to our best research system. Sensitivities were very low in order that they achieved a low false alarm rate. These systems employed by these sites were extremely complex, often involved three or four parallel branches that independently processed the data and then merged results. These systems were not particularly suited towards real-time implementation. A more thorough analysis of this will be presented at IEEE SPMB 2020, where we have a special session on this challenge.

We did observe that these systems were using artifact reduction technology. We implemented an artifact reduction system based on what we learned from these systems using an off-the-shelf open source system. Unfortunately, though this system appeared to successfully suppress isolated artifacts, it did not improve the overall system performance. That was a similar conclusion that we drew based on our own internally developed artifact reduction technology. On top of that, the off-the-shelf technology was not well-suited to a real-time implementation and added significant amounts of latency. Therefore, we terminated that effort and focused more on new deep learning architectures and low latency approaches.

Looking forward, we continue to track the deep learning literature for advances. We recently learned of the emergence of a new deep learning approach based on CNNs and transformers known as a Conformer. Conformers have produced some impressive results for speech recognition. We expect, time permitting, that we will explore this approach. We are also studying self-supervised learning approaches to see if these will allow us to leverage the vast amount of untranscribed data that we have. We think the combination of these two should deliver a significant improvement in performance.

Specific Objectives:

Our specific technical objectives, which align with the activities above, were:

(1) Performance: At the beginning of this project, our best seizure detection research system was performance at a sensitivity of approximately 30% with a false alarm rate of 7 per 24 hours. Our clinical target is a sensitivity of 75% with a false alarm rate of 1 per 24 hours. Our primary technical goal was to close this gap by exploring new deep learning architectures and leveraging more data.

(2) Real-time: The system we began the project with was understandably a research system. It had an infinite amount of latency and ran slower than real-time. Our goal was to reduce the system to rela-time on a standard 1.7 GHz processor with a GPU, and equallly important, to reduce latency to less than 30 secs so that the system could be used in crticial-care settings.

We have made significant progress on both of these though there is still a performance gap.

Significant Results:

1. The performance of our best research system now stands at 45% sensitivity with 5.7 false alarms per 24 hours. The real-time system's performance is very close to this and still undergoing some optimization. While we have not reached our performance targets, informal feedback on the performance of this system by clinicians has been very positive. Further, our annotation team, which uses this technology to triage data, has commented that the alignments of hypotheses produced by the system are significantly improved. We feel the system in its current state is not perfect, but usable in clinical settings.

2. We have also decreased the latency of our best real-time system to 11 seconds. We accomlished this by adopting a new architecture based on a pretrained network, and by overhauling the heuristic postprocessing steps. We are still working to reduce the latency to 6 seconds, but feel even a 11 seconds the system is usable.

3. As mentioned above, we also seized upon the opportunity to co-organize a comunity-wide challenge. The web site for this challenge is here: https://neureka-challenge.com/. The leaderboard posting the final results is here: https://neureka-challenge.com/results/. We have establshed an accepted industry-wide benchmark that shows our technolog is on par with the best submissions. We also learned more about the problem by studying the submissions from the leading systems.

Key outcomes or Other achievements:

When the opportunity arose to participate in an industry-wide challenge, we eagerly agreed to co-organize the event rather than compete. The main reason for this is that we wanted to make sure the competition was scored properly, using the scoring software we provide, so that performance could be directly and fairly compared with our research. A secondary reason was that we wanted to ensure the data was used properly, since the originators of the competition approached us about using our TUH EEG Seizure Detection Corpus.

Competitions such as this serve two important functions. First, they elevate the status of the research problem and get the community focused on it. Second, they establish industry-accepted benchmarks that make it much easier to publish research. Reviewers of top journals have been skeptical of the numbers we have been publishing arguing "it seems like the technology doesn't work." They have not accepted our counter-arguments that this is a difficult problem. Now we have the necessary evidence to show that our work has scientific merit.

Similarly, because of the competion, we have had about 750 downloads of our data. It raised community awareness of the problem and the resources we provide. In fact, there is a special session at the IEEE Signal Processing in Medicine and Biology Symposium (IEEE SPMB 2020) on the challenge. This is a conference we organize annually. This year, due to the challenge, a large number of papers are focused on the seizure detection problem.

*** What opportunities for training and professional development has the project provided?**

This project has employed four undergraduates in three capacities: web developer, application programmer and software engineer. The student doing web development had the opportunity to learn state of the art methods in full stack web development including using the package Boost to develop and maintain web pages. The application engineer position had the opportunity to learn a broad range of skills including real-time DSP programming, streaming interfaces, Unix shell programming and, of course, implementation of signal processing and deep learning in Python. The software engineering position involved learning how to design extensible software systems than facilitate the integration and adaptation of research software.

One graduate student was employed on the project. He has been learning how to develop low-latency real-time signal processing systems that stream data, as well as how to apply and optimize deep learning software in Python.

All students in the project also learned how to conduct a community-wide technology evaluation involving a blind data evaluation. These kinds of data challenges are quite common today and represent an important way technology is advanced. Not only do they learn how to organize and manage data sets, but they also got a chance to learn how other research groups attack the problem.

All students have also been trained on how to present their work at a professsional conference, and how to document their work. We use software management systems and a weekly reporting structure that emphasize good technical communication.

Finally, because this is a commercialization project, they have learned how to present technical material to investors and other commercial interests.

*** Have the results been disseminated to communities of interest? If so, please provide details.**

The Neureka Challenge was an open source data challenge that created a significant amount of visibility in the community. A web site was maintained that is located here: https://neureka-challenge.com/. A Slack channel was also created. Results of the competition are posted here: https://neureka-challenge.com/results/. The winning participants were required to share their code, which was done via Github sites. We did not participate as a site since we had access to the evaluation data. However, we are able to cite comparisons of performance in our upcoming publications.

Scoring software was also distributed as open source software: https://www.isip.piconepress.com/projects/tuh_eeg/downloads/nedc_eval_eeg/. We provided extensive tutorials and support to make it easy for participants to score their systems. We also adapted the file formats to make it easy for people to interface to the software.

Of course, perhaps the most significant artifact distributed was the data itself. We have been developing annotated EEG data for years and have over 3,500 subscribers of this data. It is freely available from our project web site: https://www.isip.piconepress.com/projects/tuh_eeg/html/downloads.shtml. This is the largest resource of its type in the world.

Finally, information about the project is being disseminated through publications at a conference we host annually - IEEE SPMB (2019 and 2020).

*** What do you plan to do during the next reporting period to accomplish the goals?**

We have one main challenge at this point: improve sensitivity and reduce the false alarm rate. There are two things we are exploring: (1) leveraging pretrained networks, such as ResNet, and (2) integrating more data using self-supervised training techniques. Both topics are trends in the machine learning community at present and the subject of a significant amount of research.

We have converted our primary approach to seizure detection from a sequential signal processing problem to more of an image processing approach - treating the EEG multichannel signal as an image. We believe this is closer to what neurologists do when interpreting EEGs - they clearly focus on waveshapes. Many signal processing application of machine learning seem to be moving in this direction. These approaches are exploiting pretrained networks that seem to avoid the problems associated with divergence or vanishing gradients during training. These pretrained networks can be quickly adapted to a specific task with a small amount of training data. This approach seems to be surprisingly effective for EEGs. Though it has not resulted in significant improvements in sensitivity, it has allows us to reduce delay from 30 seconds to five to ten seconds. That is significant for clinical applications requiring low delay.

We are also investigating ways to leverage the vast amount of unannotated data that we have as part of the TUH EEG Corpus. While we have manually annotated a significant amount of data, it is still a relatively small portion of the overall corpus (less than 10%). Over the past year we have investigated the use of autoencoders, which are a form of unsupervised training, to improve the front end portion of our system. We had hoped to replace features with an autoencoder, but that has not significantly improved performance. We are now looking at a new class of networks trained using what is referred to as self-supervised training. The main benefit of this approach is that it does not require detailed annotations of the data, so it creates the potential to process a much larger portion of our corpus.

Finally, with COVID-19, it has been hard to be active in the entrepreneurial aspects of this project - namely seeking out investors and potential funders of our commercialization activities. We need to spend more effort on that aspect of the program now that we have established state of the art in performance and have a better understanding of how our technology compares to others. We feel we have a strong commercialization story to tell.

Supporting Files

Filename	Description	Uploaded By	Uploaded On
fig_01.pdf	A DET curve comparing the performance of our systems.	Joseph Picone	11/08/2020
fig_02.pdf	An analysis of the latency of our channel-based LSTM system.	Joseph Picone	11/08/2020

Products

Books

Obeid, Iyad; Selesnick, Ivan; Picone, Joseph (2020). *Signal Processing in Medicine and Biology: Emerging Trends in Research and Applications 1.* Obeid, Iyad; Selesnick, Ivan; Picone, Joseph. Springer. New York, New York, USA. Status = PUBLISHED; Acknowledgment of Federal Support = No ; Peer Reviewed = Yes ; DOI: 10.1007/978-3-030-36844-9

Obeid, Iyad; Selesnick, Ivan; Picone, Joseph (2021). *Biomedical Signal Processing: Innovation and Applications 1.* 1. Obeid, Iyad; Selesnick, Ivan; Picone, Joseph. Springer. New York, New York, USA. Status = AWAITING_PUBLICATION; Acknowledgment of Federal Support = No ; Peer Reviewed = Yes ; OTHER: Not Yet Available

Book Chapters

Golmohammadi, M.; Shah, Vinit; Obeid, Iyad; Picone, Joseph (2020). Deep Learning Approaches for Automatic Analysis of EEGs. *Deep Learning: Algorithms and Applications 1st.* Obeid, I.; Picone, J.. Springer-Verlag. New York, New York, USA. 233. Status = PUBLISHED; Acknowledgement of Federal Support = No ; Peer Reviewed = Yes ; DOI: 10.1007/978-3-030-36844-9.

Shah, Vinit; Golmohammadi, Meysam; Obeid, Iyad; Picone, Joseph (2021). Objective Evaluation Metrics for Automatic Classification of EEG Events. *Biomedical Signal Processing: Innovation and Applications 1.* 1. Obeid, Iyad; Selesnick, Ivan; Picone, Joseph. Springer-Verlag. New York, New York, USA. 1. Status = ACCEPTED; Acknowledgement of Federal Support = Yes ; Peer Reviewed = Yes ; OTHER: https://www.isip.piconepress.com/publications/unpublished/book_sections/2021/springer/metrics/.

Inventions

Journals or Juried Conference Papers

View all journal publications currently available in the NSF Public Access Repository for this award.

The results in the NSF Public Access Repository will include a comprehensive listing of all journal publications recorded to date that are associated with this award.

Ferrell, Sean and Mathew, Vineetha and Refford, Matthew and Tchiong, Vincent and Ahsan, Tameem and Obeid, Iyad and Picone, Joseph. (2020). The Temple University Hospital EEG Corpus: Electrode Location and Channel Labels. *Institute for Signal and Information Processing Report*. 1 (1) 1-9. Status = Deposited in NSF-PAR doi:https://doi.org/ ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Shawki, Nabila and Elseify, Tarek and Cap, Thao and Shah, Vinit and Obeid, Iyad and Picone, Joseph. (2020). A Deep Learning-Based Real-time Seizure Detection System. *IEEE Signal Processing in Medicine and Biology Symposium SPMB*. 1 (1) 1-4. Status = Deposited in NSF-PAR doi:https://doi.org/ ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Shah, Vinit and Golmohammadi, Meysam and Obeid, Iyad and Picone, Joseph. (2021). Objective Evaluation Metrics for Automatic Classification of EEG Events. *Biomedical Signal Processing: Innovation and Applications*. 1 (1) 1-26. Status = Deposited in NSF-PAR doi:https://doi.org/ ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Jean-Paul, S. and Elseify, T. and Obeid, I. and Picone, J.. (2019). Issues in the Reproducibility of Deep Learning Results. *IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*. 1 (1) 1 to 4. Status = Deposited in NSF-PAR doi:https://doi.org/10.1109/SPMB47826.2019.9037840 ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Ochal, Domenic and Rahman, Safwanur and Ferrell, Sean and Elseify, Tarek and Obeid, Iyad and Picone, Joseph. (2020). The Temple University Hospital EEG Corpus: Annotation Guidelines. *Institute for Signal and Information Processing Report*. 1 (1) 1-28. Status = Deposited in NSF-PAR doi:https://doi.org/ ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Rahman, Safwanur Hamid. (2020). Improving the Quality of the TUSZ Corpus. *IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*. 1 (1) 1-5. Status = Deposited in NSF-PAR doi:https://doi.org/ ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Shah, Vinit and Obeid, Iyad and Picone, Joseph and Ekladious, George and Iskander, Ray and Roy, Yannick. (2020). Validation of Temporal Scoring Metrics for Automatic Seizure Detection. *Proceedings of the IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*. 1 (1) 1-5. Status = Deposited in NSF-PAR doi:https://doi.org/ ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Hamid, Ahmed and Gagliano, Katherine and Rahman, Safwanur and Tulin, Nikita and Tchiong, Vincent and Obeid, Iyad and Picone, Joseph. (2020). The Temple University Artifact Corpus: An Annotated Corpus of EEG Artifacts. *IEEE Signal Processing in Medicine and Biology Symposium SPMB*. 1 (1) 1. Status = Deposited in NSF-PAR doi:https://doi.org/ ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Rahman, S. and Miranda, M. and Obeid, I. and Picone, J.. (2019). Software and Data Resources to Advance Machine Learning Research in Electroencephalography. *IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*. 1 (1) 1 to 4. Status = Deposited in NSF-PAR doi:https://doi.org/10.1109/SPMB47826.2019.9037851 ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Golmohammadi, Meysam and Shah, Vinit and Obeid, Iyad and Picone, Joseph. (2020). Deep Learning Approaches for Automatic Seizure Detection from Scalp Electroencephalograms. *Signal Processing in Medicine and Biology: Emerging Trends in Research and Applications*. 1 (1) 233-274. Status = Deposited in NSF-PAR doi:https://doi.org/10.1007/978-3-030-36844-9 ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Kiral, Isabell and Roy, Subhrajit and Mummert, Todd and Braz, Alan and Tsay, Jason and Tang, Jianbin and Asif, Umar and Schaffter, Thomas and Mehmet, Eren and Picone, Joseph and Obeid, Iyad and Marques, Bruno De and Maetschke, Stefan and Khalaf, Rania and Rosen-Zvi, Michal and Stolovitzky, Gustavo and Mirmomeni, Mahtab and Harrer, Stefan and Yanagisawa, Hirki and Iwamori, Toshiya and Madan, Piyush and Qin, Yong and Ma, Li and Ti, Wei Lian and Liu, Wen and Mei, Jing and Hensley, Sharon and Chandra, Rachita and Hake, Paul and Henessy, Richard and Babaali, Parisa and Yuenreed, Gigi and Kather, Ryan and Arcos-Diaz, Dario and Cherner, Michael. (2019). The Deep Learning Epilepsy Detection Challenge: Design, Implementation, and Test of a New Crowd-Sourced AI Challenge Ecosystem. *Challenges in Machine Learning Competitions for All (CiML)*. 1 (1) 1-3. Status = Deposited in NSF-PAR doi:https://doi.org/ ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Lin, Rebecca and Marquez, Destiny and Jacobson, Mercedes and Castaldi, Hannah and Buckman, Samuel and Shah, Vinit and Picone, Joseph. (2020). Accuracy of Automated Machine Learning Software in Identifying EEGs with Prolonged Seizures. *Annual Meeting of the American Academy of Neurology (AAN)*. 1 (1) 1. Status = Deposited in NSF-PAR doi:https://doi.org/ ; Federal Government's License = Acknowledged. (Completed by Picone, Joseph on 10/29/2020) Full text Citation details

Shah, V., Obeid, I., Picone, J., Iskander, R., & Roy, Y. (2020). Validation of Temporal Scoring Metrics for Automated Seizure Detection. Proceedings of the IEEE Signal Processing in Medicine and Biology Symposium. https://www.isip.piconepress.com/publications/unpublished/conferences/2020/ieee_spmb/scoring/. Status = ACCEPTED.

Licenses

Other Conference Presentations / Papers

Shawki, Nabila; Elseify, Tarek; Cap, Thao; Shah, Vinit; Obeid, Iyad; Picone, Joseph (2020). *A Deep Learning-Based Real-time Seizure Detection System*. IEEE Signal Processing in Medicine and Biology Symposium (SPMB). Philadelphia, Pennsylvania, USA. Status = ACCEPTED; Acknowledgement of Federal Support = Yes

Lin, R.; Marquez, D.; Jacobson, M.; Castaldi, H.; Buckman, S.; Shah, V.; Picone, J. (2020). *Accuracy of Automated Machine Learning Software in Identifying EEGs with Prolonged Seizures*. Annual Meeting of the American Academy of Neurology (AAN). Philadelphia, Pennsylvania, USA. Status = PUBLISHED; Acknowledgement of Federal Support = No

Rahman, S. Hamid, A.; Ochal, D.; Obeid, I.; Picone, J. (2020). *Improving the Quality of the TUSZ Corpus*. Proceedings of the IEEE Signal Processing in Medicine and Biology Symposium (SPMB). Philadelphia, Pennsylvania, USA. Status = ACCEPTED; Acknowledgement of Federal Support = Yes

Jean-Paul, Shmyrde; Elseify, Tarek; Obeid, Iyad; Picone, Joseph (2019). *Issues in the Reproducibility of Deep Learning Results*. IEEE Signal Processing in Medicine and Biology Symposium (SPMB). Philadelphia, Pennsylvania, USA. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Rahman, Saf; Miranda, Matthew; Obeid, Iyad; Picone, Joseph (2019). *Software and Data Resources to Advance Machine Learning Research in Electroencephalography*. IEEE Signal Processing in Medicine and Biology Symposium (SPMB). Philadelphia, Pennsylvania, USA. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Kiral, I., Roy, S., Mummert, T., Braz, A., Tsay, J., Tang, J., Asif, U., Schaffter, T., Mehmet, E., Picone, J., Obeid, I., Marques, B., Maetschke, S., Khalaf, R., Rosen-Zvi, M., Stolovitzky, G., Mirmomeni, M., Harrer, S., Yanagisawa, H., Iwamori, T., Madan, P., Qin, Y., Ma, L., Ti, W., Liu, W., Mei, J., Hensley, S., Chandra, R., Hake, P., Henessy, R., Babaali, P., Yuenreed, G., Kather, R., Arcos-Diaz, D., Cherner, M. (2019). *The Deep Learning Epilepsy Detection Challenge: Design, Implementation, and Test of a New Crowd-Sourced AI Challenge Ecosystem*. Challenges in Machine Learning Competitions for All (CiML). Vancouver, Canada. Status = PUBLISHED; Acknowledgement of Federal Support = No

Hamid, Ahmed; Gagliano, Katherine; Rahman, Safwanur; Tulin, Nikita; Tchiong, Vincent; Obeid, Iyad Picone, Joseph (2020). *The Temple University Artifact Corpus: An Annotated Corpus of EEG Artifacts*. IEEE Signal Processing in Medicine and Biology Symposium (SPMB). Philadelphia, Pennsylvania, USA. Status = ACCEPTED; Acknowledgement of Federal Support = Yes

Other Products

Audio or Video Products.

We prepared a demonstration/advertising video available from this URL:

https://www.isip.piconepress.com/projects/nsf_pfi_tt/resources/videos/realtime_eeg_analysis/current/video_2.5.1.mp4

We use this to introduce people to the technology.

Audio or Video Products.

We prepared a video for the NSF STEM DIVE competition. The video is available from this URL:

https://www.isip.piconepress.com/projects/nsf_pfi_tt/resources/videos/stem_dive/video.mp4

Other Publications

Ochal, Domenic; Rahman, Safwanur; Ferrell, Sean; Elseify, Tarek; Obeid, Iyad; Picone, Joseph (2020). *The Temple University Hospital EEG Corpus: Annotation Guidelines*. The goal of this document is to describe the file formats used to store annotations for the Temple University Hospital EEG (TUEG) Corpus (Obeid & Picone, 2016). Subsets of the corpus have been manually annotated (Veloso et al., 2017) and are available from our project web site (Choi et al., 2017). These annotations are stored in two formats: a label file (*.lbl*) that represents an annotation as a hierarchical graph, and a time-synchronous event file (*.tse*) that represents an annotation as a flat series of events with start and stop times, type of seizure, and probability. In this document, we describe each of these formats. Tools to read and display this information are also available from our project web site (Capp et al., 2018; McHugh & Picone, 2016). https://www.isip.piconepress.com/publications/reports/2020/tuh_eeg/annotations/. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Ferrell, Sean; Mathew, Vineetha; Refford, Matthew; Tchiong, Vincent; Ahsan, Tameem; Obeid, Iyad; Picone, Joseph (2020). *The Temple University Hospital EEG Corpus: Electrode Location and Channel Labels*. The goal of this report is to describe to users of the TUH EEG Corpus four important concepts that must be understood to correctly retrieve EEG signals from a data file (e.g., an EDF file). The four key concepts described in this document are: (1) physical placement: the location of the electrodes on the scalp, (2) unipolar montage: the differential recording process used to reduce noise, (3) channel labels: the system used to describe the channels, or digital signals, represented in a computer file and (4) bipolar montages: the differential mapping used to accentuate clinically-relevant events in the signal. This report is not intended to be a primer on the electrophysiology of an EEG, which is a subject unto itself, or a tutorial on how neurologists interpret EEGs. This report simply explains how the signal data in an EEG file must be accessed to accurately support clinical applications (e.g., manual interpretation or annotation of an EEG) and research applications (e.g., automatic interpretation using machine learning). https://www.isip.piconepress.com/publications/reports/2020/tuh_eeg/electrodes/. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Patent Applications

Technologies or Techniques

Thesis/Dissertations

Websites or Other Internet Sites

NSF PFI-TT: Real-time Analysis of Electroencephalograms in an Intensive Care Environment
https://www.isip.piconepress.com/projects/nsf_pfi_tt/

We use this web site to disseminate information about the project to the community.

Participants/Organizations

Research Experience for Undergraduates (REU) funding

Form of REU funding support:	REU supplement
How many REU applications were received during this reporting period?	12
How many REU applicants were selected and agreed to participate during this reporting period?	4
REU Comments:	

What individuals have worked on the project?

Name	Most Senior Project Role	Nearest Person Month Worked
Obeid, Iyad	PD/PI	1
Picone, Joseph	Co PD/PI	1
Shah, Vinit	Graduate Student (research assistant)	6
Shawki, Nabila	Graduate Student (research assistant)	2

Name	Most Senior Project Role	Nearest Person Month Worked
Cap, Thao	Undergraduate Student	1
Elseify, Tarek	Undergraduate Student	1
Miranda, Matthew	Undergraduate Student	1
Vorwick, Lynn	Undergraduate Student	1

Full details of individuals who have worked on the project:

Iyad Obeid
Email: iobeid@temple.edu
Most Senior Project Role: PD/PI
Nearest Person Month Worked: 1

Contribution to the Project: Dr. Obeid focuses on commercialization of the technology and provides subject matter expertise.

Funding Support: None

Change in active other support: No

International Collaboration: No
International Travel: No

Joseph Picone
Email: joseph.picone@gmail.com
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: 1

Contribution to the Project: Project management and technical leader for the development of the core technology and the real-time system.

Funding Support: None.

Change in active other support: Yes

International Collaboration: No
International Travel: No

Vinit Shah

Email: tug14467@temple.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 6

Contribution to the Project: Lead algorithm designer. Artifact reduction. Deep learning architectures.

Funding Support: No additional funding. Transferred from another externally funded project for Summer 2019.

International Collaboration: No
International Travel: No

Nabila Shawki
Email: tuk02200@temple.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 2

Contribution to the Project: Real-time implementation. Feature extraction development using augmented information.

Funding Support: Worked 9 months as a department teaching assistant prior to Summer 2019.

International Collaboration: No
International Travel: No

Thao Cap
Email: tuj64267@temple.edu
Most Senior Project Role: Undergraduate Student
Nearest Person Month Worked: 1

Contribution to the Project: Software engineer. Supports the development of some of software and infrastructure.

Funding Support: Nothing other than this grant.

International Collaboration: No
International Travel: No

Tarek Elseify
Email: tug35668@temple.edu
Most Senior Project Role: Undergraduate Student
Nearest Person Month Worked: 1

Contribution to the Project: Web developer and software engineer. Tarek helped develop the user interface for the real-time demo and addressed issues with Python efficiency.

Funding Support: None

International Collaboration: No
International Travel: No

Matthew Miranda
Email: tuj66769@temple.edu
Most Senior Project Role: Undergraduate Student
Nearest Person Month Worked: 1

Contribution to the Project: Software engineer responsible for the development of the real-time system.

Funding Support: None.

International Collaboration: No
International Travel: No

Lynn Vorwick
Email: tug70217@temple.edu
Most Senior Project Role: Undergraduate Student
Nearest Person Month Worked: 1

Contribution to the Project: Software engineer responsible for design and testing of our Python code. Manages the released software and makes sure it conforms to our standards.

Funding Support: None other than this grant.

International Collaboration: No
International Travel: No

What other organizations have been involved as partners?
Nothing to report.

Were other collaborators or contacts involved? If so, please provide details.
Nothing to report

Impacts

What is the impact on the development of the principal discipline(s) of the project?

Conducting the 2020 Neureka Epilepsy Challenge was an important step because it established industry-standard benchmarks for seizure detection. These submissions were all scored with exactly the same scoring software, which is a first for this industry. Performance was much lower than anyone expected and verified that the technology we have developed is state of the art.

The latency of the real-time system is now acceptable for clinical applications. Achieving this required overhauling our software implementation and changing many of the underlying algorithms.

The use of pre-trained networks has made an impact, but not as much as we or other had anticipated. The improvements we are seeing are much smaller than what has been observed in other disciplines such as speech and image processing.

What is the impact on other disciplines?

The results we have achieved with our deep learning techniques are of interest to the general machine learning community because they are less optimistic than what has been demonstrated with other disciplines.

Obviously, a real-time low-latency system has direct impact on healthcare related fields such as neuroscience/neurology since the system can be used in clinical applications to accelerate productivity.

What is the impact on the development of human resources?

The undergraduates contributing to the project have learned a great deal about data science. They have been trained on a wide range of issues including data annotation, big data methods, Python and C++ programming, machine learning and software engineering. They learn how to design and implement large software systems in a structured environment.

The graduate students have been trained this year on how to turn high-level research implementations into real-time systems. This involved training them on some fundamentals of digital signal processing and introducing them to the process of lowering the latency of a system. Since most graduate students these days like to program in Python and don't really think about real-time implementations, this was an important step in their professional preparation for jobs in industry.

What was the impact on teaching and educational experiences?

In the process of developing this technology, we have produced a number of open source resources that are being used by the community at large. For example, our project report describing how EEG signals are organized in an EDF file is being widely disseminated. We also provide supporting software in Python and C++ that demonstrates how to read these files in a very flexible manner.

Also, obviously an open source challenge in which the winning participants were required to publish their implementations is a great contribution to the field. It helps less-experienced researchers quickly come up to speed on best practices.

What is the impact on physical resources that form infrastructure?

This is a software-based project, so there is no significant impact on physical resources. We do maintain a web site and an anonymous rsync server that we use to distribute software and data, and this resource is used quite frequently by the community.

What is the impact on institutional resources that form infrastructure?

Aside from our standard computing infrastructure, which includes a cluster that spans three proprietary networks, there is no significant impact on institutional resources. With this cluster, we are able to move data between HIPAA-protected networks and our main campus network, and this is a fairly unique capability.

What is the impact on information resources that form infrastructure?

Not particularly relevant to this project aside from the usual web resources.

What is the impact on technology transfer?

This project has a heavy technology transfer component. There were two main goals in the project - reach clinically acceptable levels of performance and produce a real-time system capable of being licensed to equipment providers.

We accomplished the second goal. We have a real-time system that has an acceptable level of latency.

We have also demonstrated that the performance of our system is comparable to the best performing systems. We continune to work to improve performance of course. However, the current system, we strongly believe, it better than what is commercially in use. We have begun discussions with equipment providers again to explore potential licensing arrangements.

Hence, we are much better positioned than at the beginning of the project to commercialize the technology and build on our STTR Phase I grant.

What is the impact on society beyond science and technology?

This project is directly relevant to the healthcare industry. We are optimistic that the system we have developed can be used successfully in the clinic, and look forward to better marketing of this technology once we clear COVID-19.

What percentage of the award's budget was spent in a foreign country?

None.

Changes/Problems

Changes in approach and reason for change

Performance of our best research system had stagnated. Also, we were having trouble translating that system to a real-time system. Therefore, we began exploring alternate approaches based on some related research in our group. We adopted a low-latency system that we felt was much easier to translate to a real-time system. Though it gave a very small degradation in performance, it reduced latency to approximately 6 secons.

Actual or Anticipated problems or delays and actions or plans to resolve them

We are behind schedule on improving the performance of the system to reach our clinical targets. We continue to closely monitor the machine learning literature to adopt promising techniques. We are also coontinuously analyzing performance and attacking the most significant error modalities.

Changes that have a significant impact on expenditures

None.

Significant changes in use or care of human subjects

None.

Significant changes in use or care of vertebrate animals

None.

Significant changes in use or care of biohazards

None.

Change in primary performance site location

None.

Special Requirements

Responses to any special reporting requirements specified in the award terms and conditions, as well as any award specific reporting requirements.