Automated Detection of Epileptic Seizures in Rodents for High-Throughput Analysis

Michael J. Reddick*, Alex H. Ruch*, Steven J. Tan*, Olivier Keunen, Yanrong Zhang, Max Wintermark
mreddick@stanford.edu, aruch@stanford.edu, sjtan@stanford.edu

Background

Epilepsy is a neurological disorder prevalent in the United States and is often associated with seizures due to the increased and uncontrollable firing of neurons in the brain. Unfortunately, one-third of patience do not respond to drug treatment, and thus surgery is required for treatment. A team of Stanford researchers and physicians led by Dr. Max Wintermark are developing an approach using Magnetic Resonance guided Focus Ultrasound (MRgFus) to enable targeted drug delivery and avoid surgery altogether.[1] To evaluate the treatment, researchers must count the number of seizures in treated rodents over hours of video. Thus, an automated, deep-learning approach to identify and count the number of seizures can save significant hours of labor.

Data Set

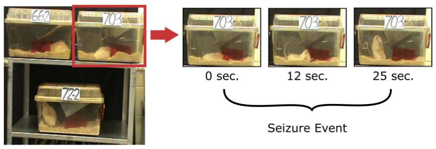

- Hundreds of hours of rat videos with set positions and cages, however only a fraction of the data in annotated
- Data is manually annotated with the start time of the seizure
- Highly asymmetric distribution of positive to negative examples – only a few ~10 sec seizures per 30 minute given video clip

Pre-processing: Encoding with Inception v4

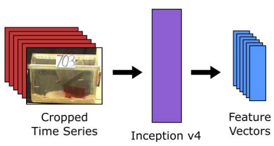

Pre-processing: Frame Differencing

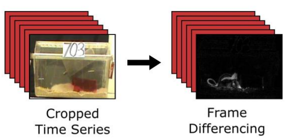

Architectures

Architecture #1: Frame Encoding with LSTM

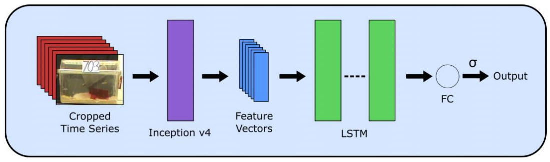

Architecture #2: Frame Differencing with 3D Conv Nets

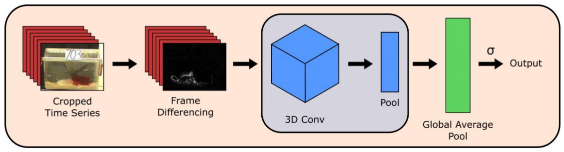

Results

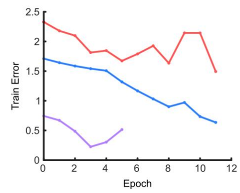

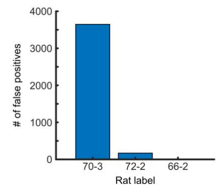

	Parameters					Metrics	
Run	RNN Size	Learning Rate	Minibatch Size	Minibatch Positive %	False Negative Penalty	Recall	F1 Score
1 (Red)	256	4.40E-05	64	12.5	62.6	100	0.064
2 (Blue)	256	2.52E-05	64	3.125	105	100	0.18
3 (Purple)	128	1.00E-04	128	3.125	18.6	0	4.40E-09

- Inception v4 + RNN produced a model with excellent recall with a sufficiently large false negative penalty and sufficiently slow learning rate.
- A large number of false positives were picked up by the model due to certain rats being more prone to seizures.
- Our model in blue performed the best with 100% recall and ~97% specificity.

Error Analysis

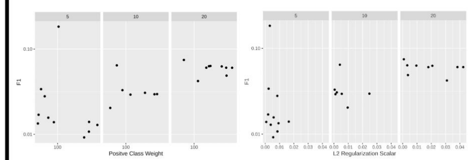

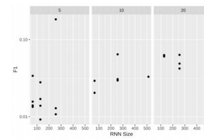

- The false negative penalty proved to be one of the most important parameters for determining model performance.
- L2 regularization parameter and the RNN size also had appreciable but less drastic affects on the model performance.

Conclusions

- The Inception + RNN architecture obtained 100% recall and 96.4% specificity in the test set.
- Large class imbalance our data contributed to very low precision.
- Most important parameters were class weight -- false negative penalty -- and the learning rate.
- An earlier layer of Inception output may better capture the features and help performance.
- Unable to fully train a network with the frame differencing followed by a 3D convolution due to time and computation limitations

References

Y. Zhang, H. Tan, E. Bertram, J.-F. Aubry, M.-B. Lopes, J. Roy, E. Dumont, M. Xie, Z. Zuo, A. Klibanov, K. Lee, M. Wintermark. "Non-invasive, focal disconnection of brain circuitry using magnetic resonance-guided low-intensity focused ultrasound to deliver a neurotoxin." Ultrasound in Medicine & Biology. 42.9 (2016): 2261-2269.
C. Szegedy, S. Ioffe, V. Vanhoucke, A. Alemi, "Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning" arXiV:602.07261 (2016).
S. Sidor, "Simple Implementation of LSTM in Tensorflow" https://gist.github.com/siemanko/b18ce332bde37e156034e5d3f60f8a23.

Acknowledgments

Thanks to Olivier Keunen, Yanrong Zhang, Max Wintermark, and the rest of the Wintermark lab for providing the videos, annotations, and scientific advice.
Thanks to Christian Aramburu for advice and ideas.

NIH National Institute of General Medical Sciences
Basic Discoveries for Better Health