

# Transfer Learning of Deep Spatiotemporal Networks to Model Arbitrarily Long Videos of Seizures

Fernando Pérez-García<sup>1,2,3</sup>[0000-0001-9090-3024], Catherine Scott<sup>4</sup>, Rachel Sparks<sup>3</sup>[0000-0003-1553-7903], Beate Diehl<sup>4</sup>[0000-0002-6896-854X], and Sébastien Ourselin<sup>3</sup>[0000-0002-5694-5340]

<sup>1</sup> Department of Medical Physics and Biomedical Engineering, University College London, UK

<sup>2</sup> Wellcome / EPSRC Centre for Interventional and Surgical Sciences (WEISS), University College London, UK

<sup>3</sup> School of Biomedical Engineering & Imaging Sciences (BMEIS), King's College London, UK

<sup>4</sup> UCL Queen Square Institute of Neurology, Department of Clinical and Experimental Epilepsy, London, UK

<sup>5</sup> National Hospital for Neurology and Neurosurgery, Department of Clinical Neurophysiology, London, UK

`fernando.perezgarcia.17@ucl.ac.uk`

**Abstract.** Detailed analysis of seizure semiology, the symptoms and signs which occur during a seizure, is critical for management of epilepsy patients. Inter-rater reliability using qualitative visual analysis is often poor for semiological features. Therefore, automatic and quantitative analysis of video-recorded seizures is needed for objective assessment. We present GESTURES, a novel architecture combining convolutional neural networks (CNNs) and recurrent neural networks (RNNs) to learn deep representations of arbitrarily long videos of epileptic seizures. We use a spatiotemporal CNN (STCNN) pre-trained on large human action recognition (HAR) datasets to extract features from short snippets ( $\approx 0.5$ s) sampled from seizure videos. We then train an RNN to learn seizure-level representations from the sequence of features. We curated a dataset of seizure videos from 68 patients and evaluated GESTURES on its ability to classify seizures into focal onset seizures (FOSs) ( $N = 106$ ) vs. focal to bilateral tonic-clonic seizures (TCSs) ( $N = 77$ ), obtaining an accuracy of 98.9% using bidirectional long short-term memory (BLSTM) units. We demonstrate that an STCNN trained on a HAR dataset can be used in combination with an RNN to accurately represent arbitrarily long videos of seizures. GESTURES can provide accurate seizure classification by modeling sequences of semiologies. The code, models and features dataset are available at <https://github.com/fepegar/gestures-miccai-2021>.

**Keywords:** Epilepsy video-telemetry · Temporal segment networks · Transfer learning.

## 1 Introduction

Epilepsy is a neurological condition characterized by abnormal brain activity that gives rise to seizures, affecting about 50 million people worldwide [8]. Seizure semiology, “the historical elicitation or observation of certain symptoms and signs” during seizures, provides context to infer epilepsy type [9]. Focal onset seizures (FOSs) start in a region of one hemisphere. If they spread to both hemispheres, they are said to *generalize*, becoming focal to bilateral tonic-clonic seizures (TCSs) [9]. In TCSs, the patient first presents semiologies associated with a FOS, such as head turning or mouth and hand automatisms. This is followed by a series of phases, in which muscles stiffen (tonic phase) and limbs jerk rapidly and rhythmically (clonic phase). TCSs put patients at risk of injury and, if the seizure does not self-terminate rapidly, can result in a medical emergency. SUDEP is the sudden and unexpected death of a patient with epilepsy, without evidence of typical causes of death. Risk of SUDEP depends on epilepsy and seizure characteristics as well as living conditions. TCSs in particular increase SUDEP risk substantially [17]. In a small number of SUDEP cases occurring in epilepsy monitoring units (EMUs), death was preceded by a TCS followed by cardiorespiratory dysfunction minutes after seizure offset [20]. Identifying semiologies related to increased risk of SUDEP to appropriately target treatment is an open research question. One limitation determining SUDEP risk factors is that inter-rater reliability based on qualitative visual analysis is poor for most semiological features (e.g., limb movement, head pose or eye gaze), especially between observers from different epilepsy centers [22]. Therefore, automatic and quantitative analysis of video-recorded seizures is needed to standardize assessment of seizure semiology across multicenter studies [3].

Early quantitative analysis studies of epileptic seizures evaluated patient motion by attaching infrared reflective markers to key points on the body or using cameras with color and depth streams [14,6,18,7]. These methods are not robust to occlusion by bed linens or clinical staff, differences in illumination and pose, or poor video quality caused by compression artifacts or details out of focus.

Neural networks can overcome these challenges by automatically learning features from the training data that are more robust to variations in the data distribution. Most related works using neural networks focus on classifying the *epilepsy type* by predicting the location of the epileptogenic zone (EZ), e.g., “temporal lobe epilepsy” vs. “extratemporal lobe epilepsy”, from short ( $\leq 2$  s) snippets extracted from videos of one or more seizures [4,2,1,16,13]. Typically, this is done as follows. First, the bed is detected in the first frame and the entire video is cropped so the field of view (FOV) is centered on the bed. During training, a convolutional neural network (CNN) is used to extract features for each frame in a sampled snippet. Then, a recurrent neural network (RNN) aggregates the features into a *snippet-level* representation and a fully-connected layer predicts the epilepsy type. Finally, a *subject-level* prediction is obtained by averaging all snippet-level predictions. This approach has several disadvantages. First, it is not robust to incorrect bed detection or changes in the FOV due to zooming or panning. Second, the order of semiologies is ignored, as the epilepsy

type is predicted from short snippets independently of their occurrence during a seizure. Moreover, patients with the same epilepsy type may present different seizure types. Finally, training neural networks from small datasets, as is often the case in clinical settings, leads to limited results.

The goal of this work is to compute *seizure-level* representations of arbitrarily long videos when a small dataset is available, which is typically the case in EMUs.

To overcome the challenge of training with small datasets, transfer learning from spatiotemporal CNNs (STCNNs) trained for human action recognition (HAR) can be used [13]. Although seizures are, strictly speaking, not actions, HAR models are expected to encode strong representations of human motion that may be relevant for seizure characterization. These methods are typically designed to classify human actions by aggregating predictions for snippets sampled from short clips ( $\approx 10$  s). Epileptic seizures, however, can last from seconds to tens of minutes [12]. A common aggregation method is to average predictions from randomly sampled snippets [5,10,21]. Averaging predictions typically works because most video datasets considered are trimmed, i.e., the same action occurs along most of the video duration. In our dataset, due to the nature of TCSs, more than half the frames are labeled as non-generalizing in 49/79 (62%) of the TCS videos. Therefore, simply averaging snippet-level predictions would result in a large number of seizures being misclassified as FOSs. Temporal segment networks (TSNs) [23] split videos of any duration into  $n$  non-overlapping segments and a consensus function aggregates features extracted from each segment. Therefore, we propose the use of TSNs to capture semiological features across the entirety of the seizure. We use an RNN as a consensus function to model the sequence of feature vectors extracted from the segments.

We present a novel neural network architecture combining TSNs and RNNs, which we denote Generalized Epileptic Seizure classification from video-Telemetry Using REcurrent convolutional neural networks (GESTURES), that provides full representations of arbitrarily long seizure videos. These representations could be used for tasks such as classification of seizure types, seizure description using natural language, or triage. To model the relevant patient motion during seizure without the need for object detection, we use a STCNN trained on large-scale HAR datasets (over 65 million videos from Instagram and 250,000 from YouTube) [10] to extract features from short snippets. Then, an RNN is used to learn a representation for the full duration of the seizure.

We chose as a proof of concept to distinguish between FOSs and TCSs, because the key distinction, if the discharge spreads across hemispheres, is only observed later in the seizure. This task demonstrates that we can train a model to take into account features across the entirety of the seizure. The main challenge, apart from the typical challenges in video-telemetry data described above, is distinguishing between TCSs and hyperkinetic FOSs, which are characterized by intense motor activity involving the extremities and trunk.

## 2 Materials and methods

### 2.1 Video acquisition

Patients were recorded using two full high-definition ( $1920 \times 1080$  pixels, 30 frames per second (FPS)) cameras installed in the EMU as part of standard clinical practice. Infrared is used for acquisition in scenes with low light intensity, such as during nighttime. The acquisition software (Micromed, Treviso, Italy) automatically resizes one of the video streams ( $800 \times 450$ ), superimposes it onto the top-left corner of the other stream and stores the montage using MPEG-2. See the supplementary materials for six examples of videos in our dataset.

### 2.2 Dataset description and ground-truth definitions

A neurophysiologist (A.A.) annotated for each seizure the following times: clinical seizure onset  $t_0$ , onset of the clonic phase  $t_G$  (TCSs only) and clinical seizure offset  $t_1$ . The annotations were confirmed using electroencephalography (EEG).

We curated a dataset comprising 141 FOSs and 77 TCSs videos from 68 epileptic patients undergoing presurgical evaluation at the National Hospital for Neurology and Neurosurgery, London, United Kingdom. To reduce the seizure class imbalance, we discarded seizures where  $t_1 - t_0 < 15$  s, as this threshold is well under the shortest reported time for TCSs [12]. After discarding short videos, there were 106 FOSs. The ‘median (min, max)’ number of seizures per patient is 2 (1, 16). The duration of FOS and TCS is 53 (16, 701) s and 93 (51, 1098) s, respectively. The total duration of the dataset is 298 minutes, 20% of which correspond to TCS phase (i.e., the time interval  $[t_G, t_1]$ ). Two patients had only FOS, 32 patients had only TCS, and 34 had seizures of both types. The ‘mean (standard deviation)’ of the percentage of the seizure duration before the appearance of generalizing semiology, i.e.,  $r = (t_G - t_0)/(t_1 - t_0)$ , is 0.56 (0.18), indicating that patients typically present generalizing semiological features in the second half of the seizure.

Let a seizure video be a sequence of  $K$  frames starting at  $t_0$ . Let the time of frame  $k \in \{0, \dots, K - 1\}$  be  $t_k = t_0 + \frac{k}{f}$ , where  $f$  is the video frame rate. We use 0 and 1 to represent FOS and TCS labels, respectively. The ground-truth label  $y_k \in \{0, 1\}$  for frame  $k$  is defined as  $y_k := 0$  if  $t_k < t_G$  and 1 otherwise, where  $t_G \rightarrow \infty$  for FOSs.

Let  $\mathbf{x} \in \mathbb{R}^{3 \times l \times h \times w}$  be a stack of frames or *snippet*, where 3 denotes the RGB channels,  $l$  is the number of frames, and  $h$  and  $w$  are the number of rows and columns in a frame, respectively. The label for a snippet starting at frame  $k$  is

$$Y_k := \begin{cases} 0 & \text{if } \frac{t_k + t_{k+l}}{2} < t_G \\ 1 & \text{otherwise} \end{cases} \quad (1)$$

### 2.3 Snippet-level classification

The probability  $\hat{Y}_k$  that a patient presents generalizing features within snippet  $\mathbf{x}_k$  starting at frame  $k$  is computed as

$$\hat{Y}_k = \Pr(Y_k = 1 \mid \mathbf{x}_k) = \mathcal{F}_{\theta_{z,x}}(\mathcal{C}_{\theta_x}(\mathbf{x}_k)) = \mathcal{F}_{\theta_{z,x}}(\mathbf{z}_k) \quad (2)$$

where  $\mathcal{C}_{\theta_x}$  is an STCNN parameterized by  $\theta_x$  that extracts features,  $\mathbf{z}_k \in \mathbb{R}^m$  is a vector of  $m$  features representing  $\mathbf{x}_k$  in a latent space, and  $\mathcal{F}_{\theta_{z,x}}$  is a fully-connected layer parameterized by  $\theta_{z,x}$  followed by a sigmoid function that maps logits to probabilities. In this work, we do not update  $\theta_x$  during training.

## 2.4 Seizure-level classification

**Temporal segment network** Let  $V = \{\mathbf{x}_k\}_{k=1}^{K-l}$  be the set of all possible snippets sampled from a seizure video. We define a sampling function  $f : (V, n, \gamma) \mapsto S$  that extracts a sequence  $S$  of  $n$  snippets by splitting  $V$  into  $n$  non-overlapping segments and randomly sampling one snippet per segment. There are two design choices: the number of segments  $n$  and the probability distribution used for sampling within a segment. If a uniform distribution is used, information from two adjacent segments might be redundant. Using the middle snippet of a segment minimizes redundancy, but reduces the proportion of data leveraged during training. We propose using a symmetric beta distribution ( $\text{Beta}(\gamma, \gamma)$ ) to model the sampling function, where  $\gamma$  controls the dispersion of the probability distribution (Fig. 1b). The set of latent snippet representations is  $Z = \{\mathcal{C}_{\theta_x}(\mathbf{x}_i)\}_{i=1}^n$ .

**Recurrent neural network** To perform a seizure-level prediction  $\hat{\mathbf{Y}}$ ,  $Z$  is aggregated as follows:

$$\hat{\mathbf{Y}} = \Pr(\mathbf{Y} = 1 \mid S) = \mathcal{F}_{\theta_{z,s}}(\mathcal{R}_{\theta_s}(Z)) = \mathcal{F}_{\theta_{z,s}}(\mathbf{z}) \quad (3)$$

where  $\mathcal{R}_{\theta_s}$  is an RNN parameterized by  $\theta_s$ ,  $\mathcal{F}_{\theta_{z,s}}$  is a fully-connected layer parameterized by  $\theta_{z,s}$  which uses a softmax function to output probabilities, and  $\mathbf{z}$  is a feature-vector representation of the entire seizure video, corresponding to the last hidden state of  $\mathcal{R}_{\theta_s}$ .

## 3 Experiments and results

All videos were preprocessed by separating the two streams into different files (replacing the small embedded view with black pixels), resampling to 15 FPS and  $320 \times 180$  pixels, and reencoding using High Efficiency Video Coding (HEVC). To avoid geometric distortions while maximizing the FOV and resolution, videos were cropped horizontally by removing 5% of the columns from each side, and padded vertically so frames were square. Snippets were resized to  $224 \times 224$  or  $112 \times 112$ , as imposed by each architecture. For realism, six video streams in which the patient was completely outside of the FOV were discarded for training but used for evaluation.

Experiments were implemented in PyTorch 1.7.0. We used a stratified 10-fold cross-validation, generated to ensure the total duration of the videos and ratio of FOSs to TCSs were similar across folds. Both views from the same video were assigned to the same fold, but videos from the same patient were not. This is because individual patients can present with both FOSs or TCSs, so data leakage

Table 1: Performance of the feature extractors. The number of parameters is shown in millions. AUC is the area under the precision-recall curve. Accuracy is computed for TCSs and FOSs, while  $F_1$ -score and AUC only for TCSs, represented by an asterisk (\*). Metrics are expressed as ‘median (interquartile range)’.

Model (frames)	Parameters	Features	Accuracy	$F_1$ -score*	AUC*
Wide R2D-50-2 (1)	66.8 M	2048	80.3 (33.2)	67.4 (30.2)	75.7 (38.4)
R2D-34 (1)	21.2 M	512	89.7 (27.7)	73.9 (23.6)	84.3 (28.7)
R(2+1)D-34 (8)	63.5 M	512	93.9 (18.3)	81.6 (16.9)	93.7 (13.4)
R(2+1)D-34 (32)	63.5 M	512	96.9 (12.9)	84.7 (13.4)	94.7 (11.9)

at the patient level is not a concern. We minimized the weighted binary cross-entropy loss to overcome dataset imbalance, using the AdamW optimizer [15]. The code is available at <https://github.com/fepegar/gestures-miccai-2021>.

For each fold, evaluation is performed using the model from the epoch with the lowest validation loss. At inference time, the network predicts probabilities for both video streams of a seizure, and these predictions are averaged. The final binary prediction is the consensus probability thresholded at 0.5. We analyzed differences in model performance using a one-tailed Mann-Whitney  $U$  test (as metrics were not normally distributed) with a significance threshold of  $\alpha = 0.05$ , and Bonferroni correction for each set of  $e$  experiments:  $\alpha_{\text{Bonf}} = \frac{\alpha}{e(e-1)}$ .

### 3.1 Evaluation of feature extractors for snippets

Despite recent advances in STCNNs for HAR, these architectures do not always outperform single-frame CNNs (SFCNNs) pre-trained on large generic datasets [11]. We assessed the ability of different feature extractors to model semiologies by training a classifier for snippet-level classification (Section 2.3).

We used two pre-trained versions of the STCNN R(2+1)D-34 [10] that take as inputs 8 frames ( $\approx 0.5$  s) or 32 frames ( $\approx 2.1$  s). Models were trained using weakly supervised learning on over 65 million Instagram videos and fully supervised learning on over 250,000 YouTube videos of human actions. We selected two pre-trained SFCNNs with 34 (R2D-34) and 50 (Wide R2D-50-2) layers, trained on ImageNet [24]. The SFCNNs were chosen so the numbers of layers (34) and parameters ( $\approx 65$  million) were similar to the STCNNs.

To ensure that all features datasets have the same number of training instances, we divided each video into segments of 32 frames. Then, we use the models to extract features from snippets of the required length (8, 32, or 1) such that all snippets are centered in the segments. The datasets of extracted feature vectors are publicly available [19]. We trained a fully-connected layer for 400 epochs on each feature set, treating views from the same video independently. We used an initial learning rate  $10^{-3}$  and mini-batches of 1024 feature vectors. We minimized a weighted binary cross-entropy loss, where the weight for TCSs was computed as the ratio of FOS frames to TCS frames.

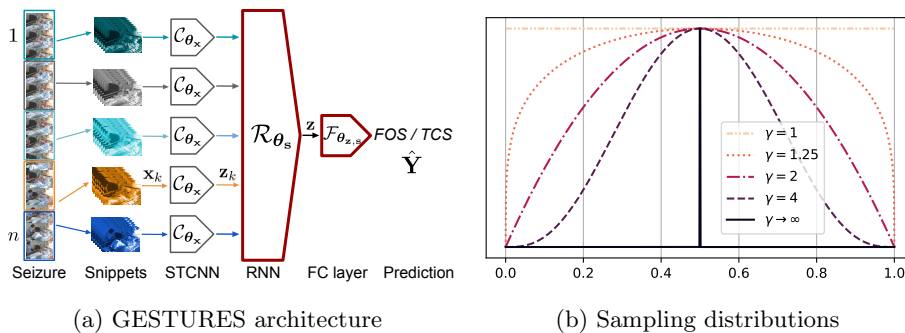


Fig. 1: Left: The GESTURES architecture. We train only the models with thick red borders. Right: Probability distributions used to sample snippets from video segments.

For evaluation, a sliding window was used to infer probabilities for all snippets. STCNNs performance was significantly better than SFCNNs ( $p < 10^{-7}$ ) (Table 1). The difference between STCNNs was not significant ( $p = 0.012$ ).

### 3.2 Aggregation for seizure-level classification

In this experiment, we compared the performance of three aggregation methods to perform seizure-level classification, using 1) the mean, 2) an RNN with 64 long short-term memory (LSTM) units and 3) an RNN with 64 bidirectional LSTM (BLSTM) units to aggregate the  $n$  feature vectors sampled from the video segments. We used the dataset of feature vectors generated by R(2+1)D-34 (8) (Section 3.1). For the task of classifying FOS and TCS, the number of segments should be selected to ensure snippets after  $t_G$ , when generalizing semiologies begin, are sampled. The theoretical minimum number of segments needed to sample snippets after  $t_G$  is  $n_{\min} = \lceil 1/(1 - r_{\max}) \rceil$ , where  $r_{\max}$  is the largest possible ratio of non-generalizing to generalizing seizure durations (Section 2.2). We can estimate  $r_{\max}$  from our dataset:  $r_{\max} = \max(r_1, \dots, r_{n_{\text{TCS}}}) = 0.93$ , where  $n_{\text{TCS}}$  is the number of TCSs, which yields  $n_{\min} = 15$  segments. We evaluated model performance using  $n \in \{2, 4, 8, 16\}$  segments per video and a sampling distribution using  $\gamma \in \{1, 1.25, 2, 4, \infty\}$ , corresponding to uniform, near semi-elliptic, parabolic, near Gaussian and Dirac’s delta distributions, respectively. For evaluation, we used  $\gamma \rightarrow \infty$ , i.e., only the central snippet of each segment. We trained using mini-batches with sequences sampled from 64 videos, and an initial learning rate of  $10^{-2}$ . We used a weighted binary cross-entropy loss for training, where the weight for TCSs was the ratio of FOSs to TCSs.

The highest accuracies were obtained using  $n = 16$  segments,  $\gamma \in \{2, 4\}$  and the BLSTM aggregator (Fig. 2). The model with the highest accuracy (98.9%) and  $F_1$ -score (98.7%) yielded 77 true positives, 104 true negatives, 2 false positives and 0 false negatives, where TCS is the positive class (Section 2.2). See

the supplementary materials for examples of videos classified correctly and incorrectly, with different levels of confidence.

## 4 Discussion and conclusion

Objective assessment of seizure semiology from videos is important to determine appropriate treatment for the diagnosed epilepsy type and help reduce SUDEP risk. Related works focus on EZ localization by averaging classifications of short snippets from multiple seizures, ignoring order of semiologies, and are not robust to variations seen in real world datasets such as changes in the FOV. Moreover, their performance is limited by the size of the training datasets, which are small due to the expense of curating datasets. Methods that take into account the sequential nature of semiologies and represent the entirety of seizures are needed.

We presented GESTURES, a method combining TSNs and RNNs to model long-range sequences of seizure semiologies. GESTURES can classify seizures into FOSs and TCSs with high accuracy. To overcome the challenge of training on limited data, we used a network pre-trained on large HAR datasets to extract relevant features from seizure videos, highlighting the importance of transfer learning in medical applications. GESTURES can take videos from multiple cameras, which makes it robust to patients being out of the FOV.

In Section 3.1 we compared STCNNs to SFCNNs for snippet-level classification. To make comparisons fair, we selected models with a similar number of layers (R2D-34) or parameters (Wide R2D-50-2). We found the larger SFCNN had worse performance, due to overfitting to the training dataset. Classification accuracy was proportional to snippet duration (Table 1), meaning that both STCNNs outperformed SFCNNs. We selected R(2+1)D-34 (8) for the aggregation experiment (Section 3.2), as performance between the two STCNNs was similar and this model is less computationally expensive.

Using LSTM or BLSTM units to aggregate features from snippets improved accuracy compared to averaging (Fig. 2), confirming that modeling the order of semiologies is important for accurate seizure representation. Model performance was proportional to the number of temporal segments, with more segments providing a denser sampling of seizure semiologies. Ensuring some dispersion in the




Fig. 2: Quantitative results for seizure-level classification. Marker brightness is proportional to the dispersion associated with the probability distribution used to sample snippets from the video segments (see Fig. 1b).



probability distributions used to sample snippets improved classification. One of the two false positives was caused by the patient being out of the FOV in one of the video streams. We did not observe overfitting to unrelated events in the videos, such as nurses in the room, to predict TCS, and models correctly discriminated between TCSs and hyperkinetic FOSs.

We demonstrated that methods designed for HAR can be adapted to learn deep representations of epileptic seizures. This enables a fast, automated and quantitative assessment of seizures. GESTURES takes arbitrarily long videos and is robust to occlusions, changes in FOV and multiple people in the room. In the future, we will investigate the potential of GESTURES to classify different types of TCSs and to localize the EZ, using datasets from multiple EMUs.

## Acknowledgments

This work is supported by the Engineering and Physical Sciences Research Council (EPSRC) [EP/R512400/1]. This work is additionally supported by the EPSRC-funded UCL Centre for Doctoral Training in Intelligent, Integrated Imaging in Healthcare (i4health) [EP/S021930/1] and the Wellcome / EPSRC Centre for Interventional and Surgical Sciences (WEISS, UCL) [203145Z/16/Z]. The data acquisition was supported by the National Institute of Neurological Disorders and Stroke [U01-NS090407].

This publication represents, in part, independent research commissioned by the Wellcome Innovator Award [218380/Z/19/Z/]. The views expressed in this publication are those of the authors and not necessarily those of the Wellcome Trust.

The weights for the 2D and 3D models were downloaded from TorchVision and <https://github.com/moabitcoin/ig65m-pytorch>, respectively.

## References

1. Ahmedt-Aristizabal, D., Nguyen, K., Denman, S., Sridharan, S., Dionisio, S., Fookes, C.: Deep Motion Analysis for Epileptic Seizure Classification. In: 2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). pp. 3578–3581 (Jul 2018). <https://doi.org/10.1109/EMBC.2018.8513031>, iSSN: 1558-4615
2. Ahmedt-Aristizabal, D., Fookes, C., Denman, S., Nguyen, K., Fernando, T., Sridharan, S., Dionisio, S.: A hierarchical multimodal system for motion analysis in patients with epilepsy. *Epilepsy & Behavior* **87**, 46–58 (Oct 2018). <https://doi.org/10.1016/j.yebeh.2018.07.028>
3. Ahmedt-Aristizabal, D., Fookes, C., Dionisio, S., Nguyen, K., Cunha, J.P.S., Sridharan, S.: Automated analysis of seizure semiology and brain electrical activity in presurgery evaluation of epilepsy: A focused survey. *Epilepsia* **58**(11), 1817–1831 (2017). <https://doi.org/https://doi.org/10.1111/epi.13907>
4. Ahmedt-Aristizabal, D., Fookes, C., Nguyen, K., Denman, S., Sridharan, S., Dionisio, S.: Deep facial analysis: A new phase I epilepsy evaluation using computer vision. *Epilepsy & Behavior* **82**, 17–24 (May 2018). <https://doi.org/10.1016/j.yebeh.2018.02.010>

5. Carreira, J., Zisserman, A.: Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4724–4733 (Jul 2017). <https://doi.org/10.1109/CVPR.2017.502>, iISSN: 1063-6919
6. Cunha, J.P.S., Vollmar, C., Li, Z., Fernandes, J., Feddersen, B., Noachtar, S.: Movement quantification during epileptic seizures: a new technical contribution to the evaluation of seizure semiology. In: Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE Cat. No.03CH37439). vol. 1, pp. 671–673 Vol.1 (Sep 2003). <https://doi.org/10.1109/IEMBS.2003.1279851>, iISSN: 1094-687X
7. Cunha, J.P.S., Choupina, H.M.P., Rocha, A.P., Fernandes, J.M., Achilles, F., Loesch, A.M., Vollmar, C., Hartl, E., Noachtar, S.: NeuroKinect: A Novel Low-Cost 3Dvideo-EEG System for Epileptic Seizure Motion Quantification. PLOS ONE **11**(1), e0145669 (Jan 2016). <https://doi.org/10.1371/journal.pone.0145669>, publisher: Public Library of Science
8. Fiest, K.M., Sauro, K.M., Wiebe, S., Patten, S.B., Kwon, C.S., Dykeman, J., Pringsheim, T., Lorenzetti, D.L., Jetté, N.: Prevalence and incidence of epilepsy. Neurology **88**(3), 296–303 (Jan 2017). <https://doi.org/10.1212/WNL.0000000000003509>
9. Fisher, R.S., Cross, J.H., French, J.A., Higurashi, N., Hirsch, E., Jansen, F.E., Lagae, L., Moshé, S.L., Peltola, J., Roulet Perez, E., Scheffer, I.E., Zuberi, S.M.: Operational classification of seizure types by the International League Against Epilepsy: Position Paper of the ILAE Commission for Classification and Terminology. Epilepsia **58**(4), 522–530 (2017). <https://doi.org/10.1111/epi.13670>
10. Ghadiyaram, D., Feiszli, M., Tran, D., Yan, X., Wang, H., Mahajan, D.: Large-scale weakly-supervised pre-training for video action recognition. arXiv:1905.00561 [cs] (May 2019), <http://arxiv.org/abs/1905.00561>, arXiv: 1905.00561
11. Hutchinson, M., Samsi, S., Arcand, W., Bestor, D., Bergeron, B., Byun, C., Houle, M., Hubbell, M., Jones, M., Kepner, J., Kirby, A., Michaleas, P., Milechin, L., Mullen, J., Prout, A., Rosa, A., Reuther, A., Yee, C., Gadepally, V.: Accuracy and Performance Comparison of Video Action Recognition Approaches. In: 2020 IEEE High Performance Extreme Computing Conference (HPEC). pp. 1–8 (Sep 2020). <https://doi.org/10.1109/HPEC43674.2020.9286249>, iISSN: 2643-1971
12. Jenssen, S., Gracely, E.J., Sperling, M.R.: How Long Do Most Seizures Last? A Systematic Comparison of Seizures Recorded in the Epilepsy Monitoring Unit. Epilepsia **47**(9), 1499–1503 (2006). <https://doi.org/https://doi.org/10.1111/j.1528-1167.2006.00622.x>
13. Karácsony, T., Loesch-Biffar, A.M., Vollmar, C., Noachtar, S., Cunha, J.P.S.: A Deep Learning Architecture for Epileptic Seizure Classification Based on Object and Action Recognition. In: ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 4117–4121 (May 2020). <https://doi.org/10.1109/ICASSP40776.2020.9054649>, iISSN: 2379-190X
14. Li, Z., Silva, A.M.d., Cunha, J.P.S.: Movement quantification in epileptic seizures: a new approach to video-EEG analysis. IEEE Transactions on Biomedical Engineering **49**(6), 565–573 (Jun 2002). <https://doi.org/10.1109/TBME.2002.1001971>, conference Name: IEEE Transactions on Biomedical Engineering
15. Loshchilov, I., Hutter, F.: Decoupled Weight Decay Regularization. arXiv:1711.05101 [cs, math] (Jan 2019), <http://arxiv.org/abs/1711.05101>, arXiv: 1711.05101
16. Maia, P., Hartl, E., Vollmar, C., Noachtar, S., Cunha, J.P.S.: Epileptic seizure classification using the NeuroMov database. In: 2019 IEEE 6th

- Portuguese Meeting on Bioengineering (ENBENG). pp. 1–4 (Feb 2019). <https://doi.org/10.1109/ENBENG.2019.8692465>
17. Nashef, L., So, E.L., Ryvlin, P., Tomson, T.: Unifying the definitions of sudden unexpected death in epilepsy. *Epilepsia* **53**(2), 227–233 (Feb 2012). <https://doi.org/10.1111/j.1528-1167.2011.03358.x>
  18. O’Dwyer, R., Silva Cunha, J.P., Vollmar, C., Mauerer, C., Feddersen, B., Burgess, R.C., Ebner, A., Noachtar, S.: Lateralizing significance of quantitative analysis of head movements before secondary generalization of seizures of patients with temporal lobe epilepsy. *Epilepsia* **48**(3), 524–530 (Mar 2007). <https://doi.org/10.1111/j.1528-1167.2006.00967.x>
  19. Pérez-García, F., Scott, C., Sparks, R., Diehl, B., Ourselin, S.: Data to support the paper ”Transfer Learning of Deep Spatiotemporal Networks to Model Arbitrarily Long Videos of Seizures” (Jul 2021). <https://doi.org/10.5522/04/14781771.v1>, publisher: University College London Type: dataset
  20. Ryvlin, P., Nashef, L., Lhatoo, S.D., Bateman, L.M., Bird, J., Bleasel, A., Boon, P., Crespel, A., Dworetzky, B.A., Høgenhaven, H., Lerche, H., Maillard, L., Malter, M.P., Marchal, C., Murthy, J.M.K., Nitsche, M., Pataraiia, E., Rabben, T., Rheims, S., Sadzot, B., Schulze-Bonhage, A., Seyal, M., So, E.L., Spitz, M., Szucs, A., Tan, M., Tao, J.X., Tomson, T.: Incidence and mechanisms of cardiorespiratory arrests in epilepsy monitoring units (MORTEMUS): a retrospective study. *The Lancet. Neurology* **12**(10), 966–977 (Oct 2013). [https://doi.org/10.1016/S1474-4422\(13\)70214-X](https://doi.org/10.1016/S1474-4422(13)70214-X)
  21. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*. pp. 568–576. NIPS’14, MIT Press, Cambridge, MA, USA (Dec 2014)
  22. Tufenkjian, K., Lüders, H.O.: Seizure Semiology: Its Value and Limitations in Localizing the Epileptogenic Zone. *Journal of Clinical Neurology* (Seoul, Korea) **8**(4), 243–250 (Dec 2012). <https://doi.org/10.3988/jcn.2012.8.4.243>
  23. Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Gool, L.V.: Temporal Segment Networks for Action Recognition in Videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**(11), 2740–2755 (Nov 2019). <https://doi.org/10.1109/TPAMI.2018.2868668>, conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence
  24. Zagoruyko, S., Komodakis, N.: Wide Residual Networks. arXiv:1605.07146 [cs] (Jun 2017), <http://arxiv.org/abs/1605.07146>, arXiv: 1605.07146