



Time-Frequency Representations of Brain Oscillations: Which One Is Better?

Harald Bârzan^{1,2}, Ana-Maria Ichim^{1,2}, Vasile Vlad Moca¹ and Raul Cristian Mureşan^{1*}

¹ Department of Theoretical and Experimental Neuroscience, Transylvanian Institute of Neuroscience, Cluj-Napoca, Romania,

² Department of Electronics, Telecommunications and Informational Technologies, Technical University of Cluj-Napoca, Cluj-Napoca, Romania

Brain oscillations are thought to subserve important functions by organizing the dynamical landscape of neural circuits. The expression of such oscillations in neural signals is usually evaluated using time-frequency representations (TFR), which resolve oscillatory processes in both time and frequency. While a vast number of methods exist to compute TFRs, there is often no objective criterion to decide which one is better. In feature-rich data, such as that recorded from the brain, sources of noise and unrelated processes abound and contaminate results. The impact of these distractor sources is especially problematic, such that TFRs that are more robust to contaminants are expected to provide more useful representations. In addition, the minutiae of the techniques themselves impart better or worse time and frequency resolutions, which also influence the usefulness of the TFRs. Here, we introduce a methodology to evaluate the “quality” of TFRs of neural signals by quantifying how much information they retain about the experimental condition during visual stimulation and recognition tasks, in mice and humans, respectively. We used machine learning to discriminate between various experimental conditions based on TFRs computed with different methods. We found that various methods provide more or less informative TFRs depending on the characteristics of the data. In general, however, more advanced techniques, such as the superlet transform, seem to provide better results for complex time-frequency landscapes, such as those extracted from electroencephalography signals. Finally, we introduce a method based on feature perturbation that is able to quantify how much time-frequency components contribute to the correct discrimination among experimental conditions. The methodology introduced in the present study may be extended to other analyses of neural data, enabling the discovery of data features that are modulated by the experimental manipulation.

Keywords: neural oscillations, time-frequency representation, machine learning, explainable AI, neurophysiology, electroencephalography

OPEN ACCESS

Edited by:

Andreas Bahmer,
University Hospital Frankfurt,
Germany

Reviewed by:

Cem Uran,
Max Planck Society, Germany
Keyvan Mahjoory,
Max Planck Institute for Empirical
Aesthetics, Germany

*Correspondence:

Raul Cristian Mureşan
muresan@tins.ro

Received: 08 February 2022

Accepted: 21 March 2022

Published: 14 April 2022

Citation:

Bârzan H, Ichim A-M, Moca VV
and Mureşan RC (2022)
Time-Frequency Representations
of Brain Oscillations: Which One Is
Better?
Front. Neuroinform. 16:871904.
doi: 10.3389/fninf.2022.871904

INTRODUCTION

Under normal and pathological conditions, brain circuits engage in rhythmic modulations of activity called neural oscillations (Buzsaki, 2006). Their frequencies span orders of magnitude, from the slow and infra-slow ones, with periods of minutes, tens of minutes, or hours (Aladjalova, 1957; Hughes et al., 2011; Belle and Diekman, 2018; Chrobok et al., 2021), to the medium ones, including

theta (4–8 Hz) or alpha (8–12 Hz), the fast ones, like beta (12–30 Hz) or gamma (30–80 Hz) (Mureşan et al., 2008), and, finally, the ultra-fast ones (>100 Hz) (Csicsvari et al., 1999; Hughes, 2008; Gourévitch et al., 2020). These different frequency bands are associated to different brain states. While some specific frequencies also emerge in pathological conditions, like Parkinson's disease (Foffani et al., 2003), it is an established fact that oscillations are notably expressed in healthy subjects. For example, slower frequencies are associated with sleep (Steriade and Amzica, 1998; Steriade, 2003), while the higher frequencies seem to support cognition and perception (Keil et al., 1999; Klimesch, 1999; Melloni et al., 2007; Başar, 2013; Barraza et al., 2020). In the mid-range, theta oscillations are pervasive in the hippocampus (Wilson and McNaughton, 1994; Skaggs et al., 1996; Buzsáki, 2002; Lega et al., 2012), while alpha rhythms seem to block sensory inputs at cortical level (Traub et al., 2020). To understand their role in normal or abnormal brain function, these ubiquitous rhythmic processes need to be precisely and reliably quantified in brain signals.

One property of neural oscillations renders their estimation in brain signals particularly difficult. Unlike several natural systems, where oscillatory processes may be sustained over longer periods of time (Stuiver and Brazianus, 1989; Huisman and Weissing, 1999; Arnaut and Ibáñez, 2020), brain oscillations are typically transient: They emerge in bursts, or packets (Tal et al., 2020; Moca et al., 2021), which often span only a few cycles in duration. To evaluate their presence in the signal, such packets need to be localized in both time and frequency, simultaneously, and this is very difficult due to the Heisenberg–Gabor uncertainty principle (Gabor, 1946; Heisenberg, 1985).

The time-frequency localization of a signal is most frequently estimated by using time-frequency representations (TFRs) (Boashash, 2015). TFRs are a very common tool used to hunt for oscillation bursts in neural data. Used extensively in the analysis of electroencephalography (EEG) (Başar et al., 2001; Herrmann and Demiralp, 2005) and local-field potential (LFP) (Wilson and McNaughton, 1994; Skaggs et al., 1996; Buzsáki, 2002) data, TFRs are computed using a plethora of spectral techniques, such as short-time Fourier transforms (STFT), Wigner-Ville distributions (WVD), wavelets, and so on. Depending on the particular methodology, different aspects may be emphasized in the signal, e.g., better temporal resolution or better frequency localization (Moca et al., 2021). In addition, some techniques are less robust to noise, which acts as a distractor, hindering the observation of finite oscillation packets in the data (Moca et al., 2021). Other techniques suffer from the smearing of the packet's representation in time, from its leakage across frequencies, banding, or so-called cross-terms, which are artifactual components introduced by the method itself (Boashash, 2015). Thus, the particular details of the method can vastly impact the properties of the estimated TFR. Furthermore, even for the same method, one may obtain TFRs with different characteristics. Indeed, some methods have one or more parameters (e.g., window size, number of cycles of the mother wavelet, order of the superlet, etc.) and their choice determines what

properties of the data are emphasized or suppressed in the representation.

One of the major issues is that it is not always clear what the “correct” (or best suited) TFR of a signal should be. Each estimation technique yields TFRs that may differ in terms of time or frequency resolution (or both), and is sensitive to some properties of the data (Moca et al., 2021). Here we propose one possibility to more objectively assess the quality of a representation: To determine how much useful information the latter extracts from neural signals.

Unfortunately, determining the information content of a TFR is never an easy task, especially when dealing with feature-rich data such as neural time-series data (e.g., EEG or LFP). This is made even harder because the items humans most readily perceive in various signal representations can display features that may not necessarily be systematically related to the experimental condition. Luckily, tools developed in the past decades in Artificial Intelligence (AI), like machine learning (LeCun et al., 2015) may come to the rescue. Such tools can be helpful to judge how much task- or stimulus-related (class-related) information there is in a given labeled dataset by measuring the prediction accuracy (% of correctly classified samples) of trained classification models on test data (Ciuparu et al., 2020). Here, we show that although the “true” time-frequency representation of a signal is not necessarily known, the “quality” of a certain TFR may be judged with the help of machine learning. We use this tool to determine how informative a TFR is about the underlying experimental conditions. We also show that explainable AI techniques (Došilović et al., 2018; Barredo Arrieta et al., 2020; Linardatos et al., 2020) can be used not only to judge the “quality” of a TFR, but also to discover where in the TFR one can find the informative oscillation packets which enable the discrimination among conditions.

MATERIALS AND METHODS

In vivo Recordings

Visual stimuli consisted of either moving bars (protocol “RF”; **Figure 2A**, top) or drifting sinusoidal gratings (0.11 cycles/deg, at a speed of 1.75 cycles/s, protocol “SRCS”; **Figure 3A**, top) of 4 different orientations and 8 moving directions in steps of 45°, each presented monocularly 10 times on an LCD monitor (Beetronics) with a resolution of 1,440 × 900 pixels and a refresh rate of 60 Hz.

Adult wild type C57/BL6J mice were anesthetized with isoflurane (5% for induction, 2–2.5% for surgery) and Xylocaine was used as local analgesic. The animals were placed in a stereotaxic frame (Stoelting Co., IL, United States) and the body temperature was maintained at 37°C using a heating pad (Harvard Apparatus). To avoid dehydration during the experiment, saline was infused intraperitoneally (IP). Silicon oil (Sigma-Aldrich) was applied to both eyes to prevent corneal drying and damage. A 2 mm circular craniotomy was performed over the left visual cortex of the animal (2–2.5 mm lateral from midline, 0–0.5 anterior to lambda) to allow insertion of the neural probe. The dura matter was left intact.

Extracellular recordings were acquired using a 32-channel silicon probe (Cambridge NeuroTech) at a sampling frequency of 32 k Samples/s (Multi Channel Systems GmbH) and subsequently band-pass filtered offline (0.1–300 Hz) using a 3rd order Butterworth infinite-impulse response (IIR) filter, applied bidirectionally. All the recordings were downsampled to 1,000 samples/s and a series of notch filters (3rd order Butterworth IIRs, bidirectional) were applied to remove the line noise and its harmonics (50, 100, and 150 Hz). We thus obtained the local-field potentials (LFP) recorded intracortically (Einevoll et al., 2013; Teleńczuk et al., 2017). To minimize animal use, multiple recordings were collected over a period of 6–8 h from each animal.

The data used throughout this study was obtained from three different animals: M048—“SRCS-1,” M60—“SRCS-2,” and M065—“RF.” The channels considered for each analysis were electrodes 19, 11, and 7, respectively.

Electroencephalography

High density EEG data was digitized at 1,024 samples/second with a 128 electrode Biosemi ActiveTwo machine during a visual recognition task. At the beginning of each trial the subjects were instructed to fixate the center of the screen (baseline) for 1.5–2 s, after which the stimulus was presented on the screen. During the stimulation period subjects were allowed to freely explore the visual scene in order to reach a decision about the identity of the stimulus (object). The stimulation period ended with a key press to signal one of the three possible outcomes: the stimulus was identified (seen), something was visible but the subject was unable to identify it (uncertain), or that no discernable shape was presented on the screen (nothing). A set of 30 shapes comprising familiar objects, fruits, and animals were processed with the “Dots” method (Moca et al., 2011). The trials were organized in 7 blocks with successively increasing visibility to a total of 210 trials per recording session. Visual stimuli (Figure 5A left) spanning 8.7° and 5.6° of visual horizontal and vertical angle, respectively, were presented on a fast response (2 ms) 22-inch LCD monitor (Samsung SyncMaster 226BW). The dots method and the experimental protocols are described in more details elsewhere (Ciuparu and Mureşan, 2016; Bârzan et al., 2021; Moca et al., 2021, 2011). Here, we used data from a single subject, as a testbed to exemplify the TFR evaluation technique.

The EEG data, sampled at 1,024 samples/s, was processed offline with Fieldtrip (Oostenveld et al., 2010). First, data was band-pass filtered to 1–150 Hz and two notch filters were also applied to reject the mains noise (50 Hz) and its first harmonic at 100 Hz. Next, channels and trials with abnormally large amplitudes were rejected manually before ICA decomposition. Finally, eye, muscle, and other artifacts were identified and removed from the ICA components and the cleaned EEG was recomposed.

Short-Time Fourier Transform

The Short-Time Fourier Transform (STFT), also commonly used to compute the *spectrogram*, is one of the most established time-frequency representation methods. The concept behind the STFT

is simple and entails computing multiple windowed Fourier transforms and arranging the resulting spectra out in time:

$$STFT_x(t, f) = \int_{-\infty}^{\infty} x(\tau) h(\tau - t) e^{-j2\pi f \tau} d\tau$$

where, x is the signal of interest and h is a time-localizing window function.

Both the time and frequency resolution of the STFT are determined by the length of the analysis window, h . A long window favors better concentration in frequency but worse temporal localization, while a short window provides good temporal localization but poor frequency resolution. Therefore, the size of the analysis window is typically set to the time scale of whatever process is being followed in the analysis, but, importantly, this timescale is unique for a given TFR (Moca et al., 2021). Thus, one cannot simultaneously and optimally track multiple timescales using this representation—this is why multiresolution methods are required, such as the continuous wavelet transform or the superlet transform, both described below.

Continuous Wavelet Transform

The Continuous Wavelet Transform (CWT) is a multiresolution technique that involves convolving the input signal with *wavelets*, which can be thought of as band-pass filter kernels whose scale is adjusted for whatever frequency is being measured. Here, we used the CWT with the complex Morlet wavelets $\psi_{f,c}$ with a known bandwidth which is defined in terms of the number of cycles c of the mother wavelet (Moca et al., 2021):

$$\psi_{f,c}(t) = \frac{1}{B_c \sqrt{2\pi}} e^{-\frac{t^2}{2B_c^2}} e^{j2\pi f t}$$

where, t is time, f is frequency, and c is the number of cycles of the mother wavelet, and

$$B_c = \frac{c}{5f}$$

By scaling the wavelet in time, therefore changing its frequency response, one can measure an arbitrary range of frequencies *via* convolution with the input signal:

$$CWT_x(t, f, c) = 2 \left| (\psi_{f,c} * x)(t) \right|^2$$

Because each wavelet scales with the frequency it measures, the time resolution of this method (relative to the timescale of processes evolving on different frequencies) is optimal. The CWT suffers from the opposite drawback, however—its frequency resolution degrades as the measured frequency is increased. To improve it, one must increase the number of cycles c , but doing so incurs the cost of worse temporal resolution.

Superlet Transform

The Superlet Transform (SLT) is a wavelet-based time-frequency estimation technique (Bârzan et al., 2021; Moca et al., 2021)

introduced to get around the degrading frequency resolution of the CWT for higher frequencies (Moca et al., 2021). By combining multiple representations in a minimum-mean cross entropic sense (by using the geometric average), the Superlet achieves time-frequency super-resolution. The SLT is defined simply as:

$$SLT_{x,o,c_1}(t,f) = \left[\prod_{i=1}^o CWT(t,f,c_i) \right]^{\frac{1}{o}}$$

where, x is the signal, o is the order of the superlet, and c_1 is the number of cycles of the base wavelet.

The intuitive explanation of the SLT is that wavelets with small temporal bandwidths have very small (close to zero) responses (coefficients) outside of a process' time window, while larger wavelets will have very small responses outside of the process' frequency characteristic, thereby the combination of short and wide wavelets is isolating the process in both time and frequency.

Wigner-Ville and Reduced Interference Distributions

Another technique widely employed in signal processing is the Wigner-Ville distribution (WVD). A full explanation regarding the WVD is outside the scope of this paper, but can be found in Debnath (2002). In short, it is the Fourier transform, expressed in time-frequency coordinates (t, f), of the signal's instantaneous autocorrelation function (IAF), expressed in time-lag coordinates (t, τ). For a single-component signal (with a single isolated feature in the time-frequency plane), such as a Gaussian atom (combination of a Gaussian with a complex sine wave), the WVD achieves the best possible time-frequency resolution. For multi-component signals (multiple features in time and/or frequency), however, cross-terms arise between the different components visible as ripples in the time-frequency field. Attempts to eliminate these cross-terms have led to the development of Reduced Interference Distributions (RIDs), which use an extra step to apply a smoothing kernel to the representation (they are called ambiguity kernels as they are applied in the ambiguity domain) (Choi and Williams, 1989; Cordero et al., 2018). Unfortunately, the smoothing kernels degrade the resolution of RIDs when compared to WVD.

$$WVD_x(t,f) = \int x\left(t + \frac{\tau}{2}\right) x\left(t - \frac{\tau}{2}\right) e^{-j2\pi f\tau} d\tau$$

Here, we use the Choi-Williams (ChoiW) distribution (Choi and Williams, 1989), a form of RID, to create time-frequency representations of neural data. It has a single additional parameter, σ , which adjusts the smoothing factor of the ambiguity kernel.

Feature Vector Extraction

Time-frequency representations were computed on a single channel for each trial (8 conditions with 10 trials each for both SRCS and RF protocols, resulting in 80 representations). The parameters for the techniques were loosely set to obtain good representations for the alpha, beta, and gamma frequency

bands: (i) STFTs using 250 ms Blackman windows, zero padded to a Fourier window of 2 s and a time step of 1 ms, (ii) CWTs with 3 cycles for the mother wavelet, (iii) SLTs with either a shorter temporal footprint ($c_1 = 2, o = 5$) or a better frequency concentration ($c_1 = 3, o = 10$), and (iv) ChoiW with $\sigma = 3$. The electrophysiological (intracranial and EEG) data was analyzed differently, according to the peculiarities of their respective experiments.

For intracranial recordings, the trial length and frequency of interest regions differ depending on the experimental protocol and the frequency characteristics of the responses. For trials with the SRCS paradigm we extracted TFRs spanning 4 s, while for RF it we considered a temporal window of 3 s, matching the difference in trial size for these different stimulation protocols. Between the three mice that were used in this analysis, each had their stimulus response in another frequency range within the gamma band, which we could normalize to about 20 Hz of bandwidth (RF: 50–70 Hz, SRCS-1: 35–55 Hz, SRCS-2: 30–50 Hz).

For EEG recordings, the trial window was chosen to be a 700 ms window (Figure 1A), occurring 300 ms after stimulus presentation to avoid stimulus onset effects. The frequency range was chosen between 5 and 55 Hz (Figure 1B).

Each analysis technique produces time-frequency representations for the given time segment and frequency range, but each of them have their unique peculiarities in achieving this, especially in the digital domain where representations are discrete. For example, the number of bins per Hz in the SLT and CWT is specified by the user, but the number of bins per second will be equal to the sampling rate of the signal, which differs between the intracranial recordings (1,000 samples/s) and the EEG recordings (1,024 samples/s). STFTs can have varying number of bins for both time and frequency axes (as they can be adjusted using the step size and window length, respectively), and the ChoiW's binning is determined (fixed) by the sampling rate. Considering the diverse range of time and frequency binning of various methods, a standardization was necessary to compare the representations. Therefore, all TFRs used here were resampled from their native bin sizes (Figure 1C) to a single bin size using bicubic interpolation (Figure 1D), a method widely used in image processing. After this resampling step all TFRs have the same time- and frequency ranges and bin sizes (1 ms and 0.5 Hz per time and frequency bin, respectively, unless otherwise specified). The last step involves compacting the obtained spectrum to a more palatable number of features for classification. This entails establishing numbers of features to obtain per second and per Hz. It essentially means downsampling the time-frequency spectra by splitting in tiles and computing the average of each tile (see Figure 1D). The spectra are then linearized and labeled to prepare them for classification (Figure 1E).

Information Content Estimation Using Machine Learning

Artificial neural networks (ANNs) were used to quantify the information content of the TFRs. The networks used here

were multilayer perceptrons (MLPs) with 2 hidden layers, each layer's size being 20 and 10% of the size of the input layer, respectively (**Figure 1E**). This network architecture was determined empirically by evaluating multiple different configurations and selecting the simplest one that had a good performance. Importantly, further optimization may be possible for specific neural datasets, e.g., by increasing the depth of the network or the sizes of the hidden layers. Here we refrained from overoptimizing the networks for the different datasets in order to provide a simple and general proof of concept. Importantly however, we found that, for the data presented here, results were robust across a wide range of architectures.

The networks used a novel activation function called Soft++, discussed in Ciuparu et al. (2020). To train the networks, we used the iRPROP+ (Igel and Hüsken, 2000), an improved version of the resilient backpropagation algorithm (Rprop). Because neuroscience datasets typically yield a large number of features but a small number of samples and, in addition, the neural networks have a large number of parameters, this renders such setups prone to overfitting. To deal with this we introduced a mild dropout strategy (Hinton et al., 2012). The full parameters of the ANNs and of the learning algorithm are given in **Table 1**.

To train and test the networks we used datasets consisting of downsampled and linearized TFRs (feature vectors) computed per experimental trial (see **Figure 1**). For the EEG dataset, this yielded a total of 150 machine learning samples (feature vectors; 60 “nothing” trials and 90 “seen” trials), while for LFP datasets it provided 80 machine learning samples (10 trials for each direction—see above). For each dataset and each run of training/testing a network, the set of machine learning samples was split randomly into 60% train and 40% validation (test) samples. Training was only performed using the samples in the training set. Furthermore, training was performed for 1,000 epochs and the best network state during the training process was saved. Due to the small number of samples in our datasets, biased sampling during splitting becomes a problem that needs to be

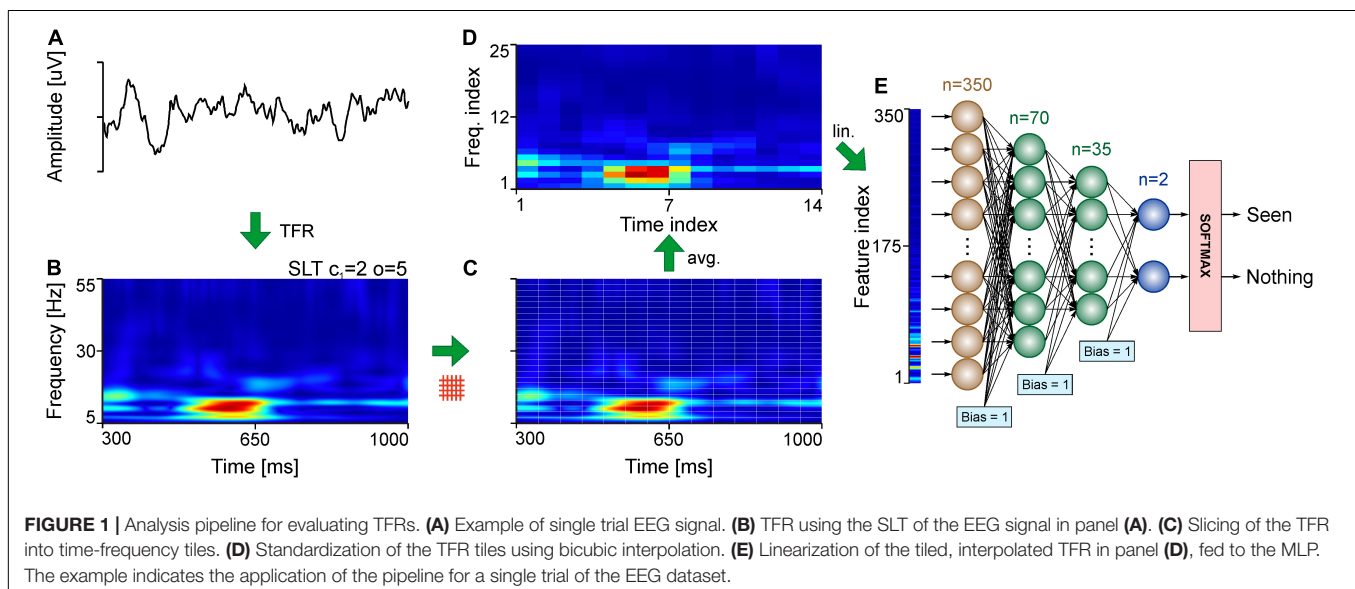
mitigated. To average out effects due to the random splitting, as well as due to the random weight initialization of the networks, all dataset splitting, training, and testing procedures were repeated 100 times for each TFR type (4 in total). Finally, to evaluate chance level (especially for unbalanced datasets), we performed the same splitting/training/testing procedure 100 times and for each TFR type while simultaneously randomizing the labels of the datasets for each run.

Feature Permutation

To evaluate how much a certain feature or set of features contributes to the classification, we developed a methodology based on feature permutation (Fisher et al., 2019). The idea behind this is to perturb each feature at the input of the ANN and to evaluate how much certain performance metrics degrade (e.g., decrease in accuracy or increase in mean-squared error). This is achieved by permuting the values of that feature between different samples such that the values of the feature no longer match the values of the other features nor the classes they belongs to. As a result, the relation between the values of the feature and class labels becomes a random one. If that feature is important for classification, then it is expected that its permutation leads to a degradation of the performance of the classifier. In addition, the importance of the feature can be quantified by measuring the amount of degradation in the performance of the classifier.

One important problem of this strategy is that the perturbation of a feature may not affect the performance of the classifier if other features are strongly correlated to that feature. Indeed, a set of strongly correlated features provides redundant information, resilient to the disturbance of some of the features in the set. Thus, the set of correlated features has to be considered as an aggregate and perturbed together.

Considering these important aspects, we devised a feature permutation methodology consisting of the following steps. The classifier is first trained using the training set and its performance on the validation set is evaluated. This is the



reference performance on the “intact” feature space. Second, the pairwise correlation (Pearson) (Nikolić et al., 2012) is computed between all the features values on the validation set. Third, features are considered one by one for the purpose of applying the permutation strategy. Before permuting the current feature, we determine the set of its correlated features by applying a threshold. Here we considered that features are part of the same set if their correlation exceeded 0.5. Next, instead of permuting only the current feature, permutation is applied to all features in the set. After the permutation, the performance of the classifier is evaluated on the new (perturbed) validation set and compared to the reference performance on the intact validation set. The permutation for each correlated feature set is performed multiple times (here we used 10 permutations), and the average change in performance is computed. Finally, the next feature is considered, its correlated set is computed, and the procedure is repeated until all features have been visited.

The permutation of a certain feature set should lead to a decrease in performance (decrease in accuracy, increase in error) but in some cases it can also lead to an increase in performance. In such a case, the feature set acts as a “distractor” for the classifier. The classification performance would be higher if the distractor feature set would be discarded. Thus, the methodology described here can be used to identify both features that are important for classification and features that act as distractors.

RESULTS

We next tested the TFR evaluation methodology on both intracranial LFP data recorded in mice and on EEG data recorded in humans. First, we focused on a receptive-field (RF) mapping protocol, where drifting, oriented bars are used to map out the structure of RFs in primary visual cortex neurons (**Figure 2A**). In this case, the passing of the bars through the RF of the cluster of neurons located at the position of the recording electrode induces vigorous gamma bursting at ~ 55 Hz. Importantly, due to the retinotopic location of the cluster of neurons relative to the drifting bar, the bursting appears at different moments throughout the trial. This temporal separation makes it such that the ANN will “perceive” increased inputs on different subsets of neurons in the input layer—as a result, classifiers learn the relationship between the TFRs and the experimental condition very robustly (**Figure 2B**; performance close to 100%) and in very few epochs (**Figure 2C**). While the SLT and STFT seem to have a slight advantage, performance is almost saturated for most TFRs. In other words, all representations contain sufficient information to unambiguously identify the direction of the moving bars. For such simple cases it is difficult to say which representation is more useful from an information content perspective—but other criteria may be applied, like the temporal and frequency localization of the gamma burst.

In a second batch of tests, we analyzed data recorded from mouse primary visual cortex when using drifting grating stimuli instead of bars. These stimuli induce periodic gamma bursting that occurs at different latencies for different conditions, as shown by all TFRs (**Figure 3A**). For this case, the STFT and

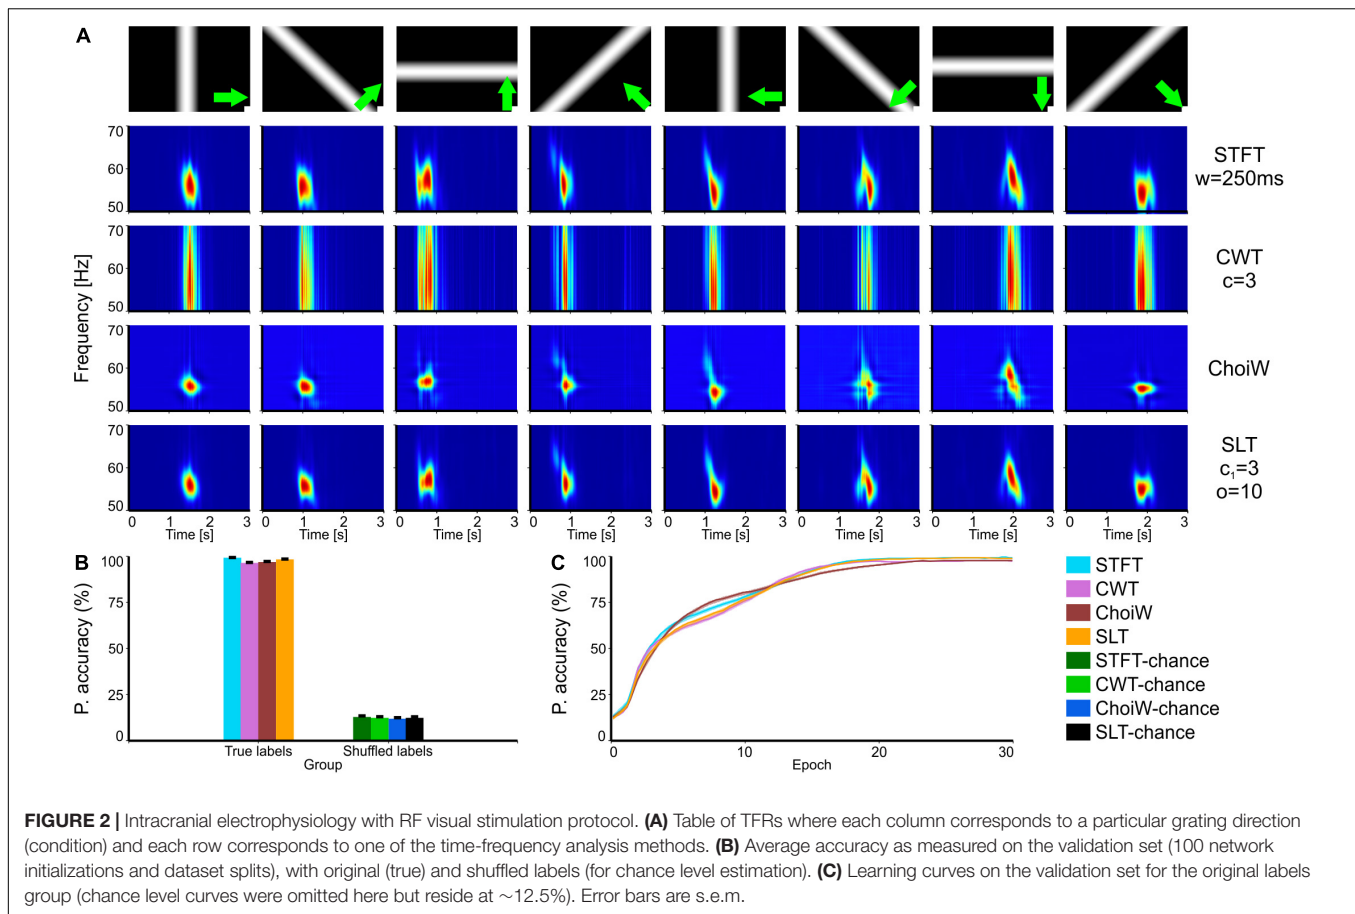
SLT provide the most informative representations (**Figure 3B**) and also enable the fastest learning (**Figure 3C**), followed by the CWT. Interestingly, the ChoiW representation, which appears to subjectively generate the sharpest frequency-resolved representation, has the lowest information content about stimulus identity. This occurs because cross-terms expressed in between the gamma bursts appear to be detrimental to the representation. Indeed, the temporal location of the bursts is important for identifying the orientation of the drifting gratings and the cross-terms of the ChoiW representation interfere with the temporal localization of the bursts.

In a third set of tests, we evaluated the information content of TFRs on data recorded with the same type of stimuli (gratings) but that exhibits more pronounced noise, evident as broadband transients in the TFRs (**Figure 4A**). In this case, the STFT provides the most informative TFR, probably because its poorer temporal resolution. Indeed, at the other end the CWT provides the least information about the stimuli because it emphasizes better the short-lived broadband noise transients (**Figure 4B**). In this case, the SLT and ChoiW representations offer a tradeoff, although learning is faster in the case of the SLT (**Figure 4C**). Thus, methods which are more sensitive to brief transients, like the CWT, provide representations that are less accurate when broadband noise contaminates the spectrum.

Intracranial electrophysiology data, such as the LFP, exhibits in general cleaner spectra than extracranial data, such as the EEG. The reason is that in extracranial signals sources of noise are more prevalent. In addition, data in **Figures 2–4** was recorded in anesthetized animals, where cortical dynamics may also be more predictable, and the experiment better controlled in terms of visual input. To have a harder and more realistic testbed, we used EEG data recorded in adult human volunteers. We evaluated how much information various TFRs contain about the conscious perception of one human subject on images of dots representing objects with various degrees of visibility (Moca et al., 2011). Classifiers were thus trained to predict the response of the subject (object not seen, i.e., “nothing,” or object “seen”) (Moca et al., 2011) by using exclusively the TFRs computed on the EEG signals recorded from the occipital electrode (Oz; **Figure 5A**).

As expected, the case of the EEG data is much more difficult for classifiers, which attain a top average performance around 70% (**Figure 5B**; chance level is $\sim 60\%$, given the class unbalance—90 trials with “seen” response compared to 60 trials with “nothing” response). In this case, the best performance was attained on the SLT representation, where classifiers also learned the fastest (**Figure 5C**), closely followed by the ChoiW distribution. The worst performer was the STFT, which provided the lowest accuracy and slowest learning curve. Thus, for complex spectra, the more advanced methods provide representations with a net advantage over the classically used TFRs, such as the STFT and the CWT.

The results presented so far indicate that various TFRs provide an informative spectral landscape that is correlated to various degrees with the experimental condition. However, it is not sufficient to determine that a representation is informative—ideally one needs to also find out which components of that representation are important and to what degree. To this end,



we developed a methodology that is able to quantify how much certain time-frequency components contribute to the accuracy of the classifiers (see section “Materials and Methods”). The methodology is based on feature permutation (Fisher et al., 2019), an algorithm belonging to the larger class of explainable AI techniques (Došilović et al., 2018; Barredo Arrieta et al., 2020; Linardatos et al., 2020).

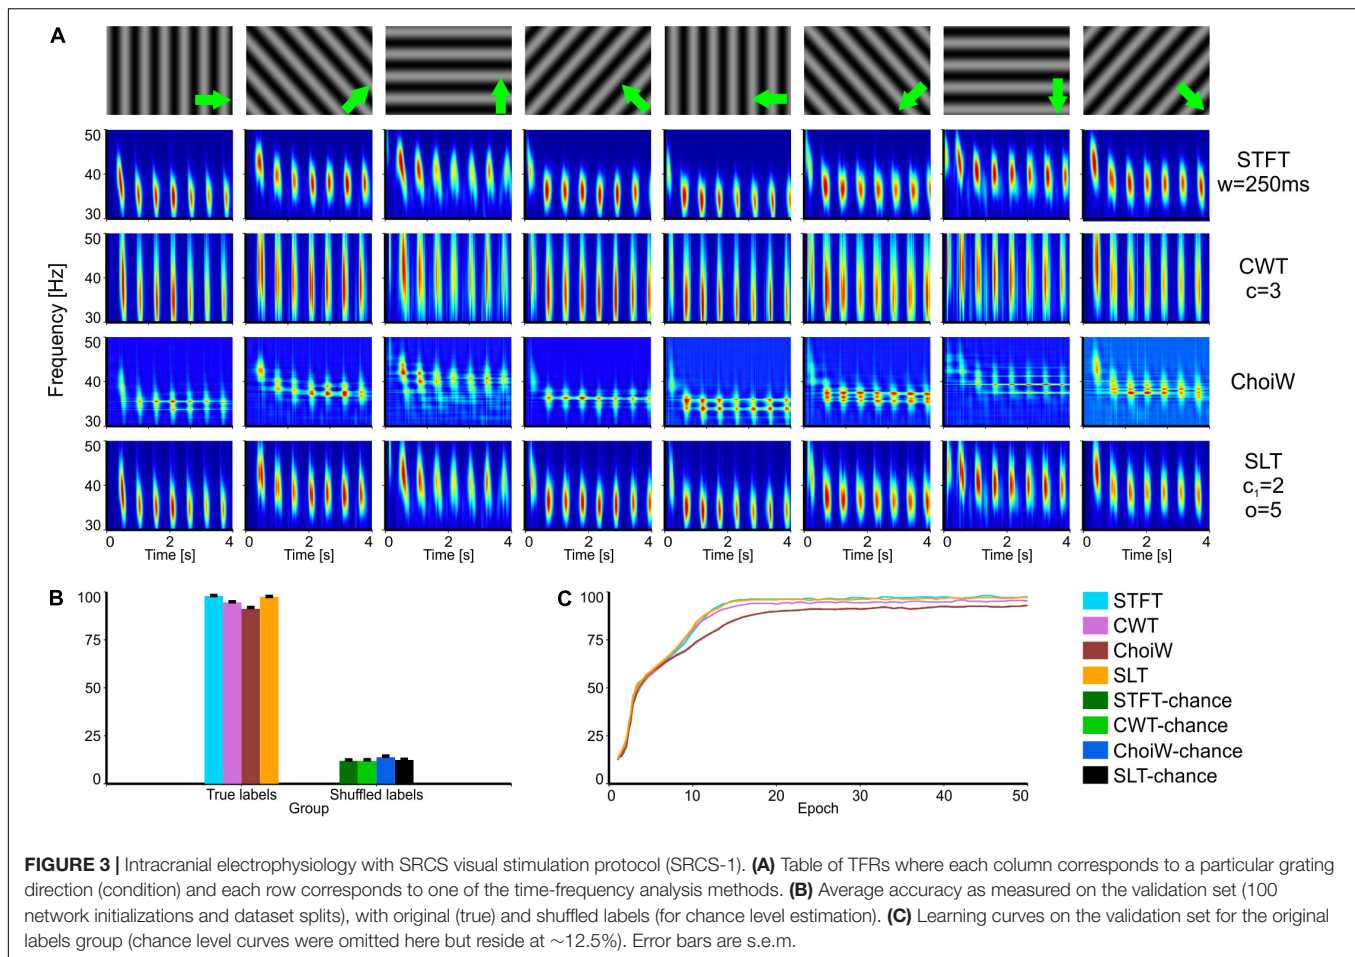
To determine what TFR components enable the differentiation of perceptual conditions in the EEG data from **Figure 5** (“seen” vs. “nothing” responses), we first computed the TFRs for all the trials and created a fixed random split into train and validation sets for all types of TFR. This ensured that a source of variability was eliminated and that all models were trained and tested on the exact same sets of experimental data.

Using the fixed training set, we trained multiple classifiers on the TFRs computed with different methods and selected the best networks, reaching the following top accuracy on the validation set: SLT: 80%, ChoiW: 78.3%, CWT: 71.7%, and STFT: 70%. Next, we applied the feature permutation technique on sets of correlated features (see section “Materials and Methods”) and computed the increase in mean-squared error (MSE) and the amount of decrease in prediction accuracy. These two quantities provide overlapping, but not identical information. If the increase in MSE is a more “continuous” measure of feature importance, the decrease in prediction accuracy is more non-linear (for

example, the degradation of a feature may increase error but not sufficiently to generate a misclassification).

Before permuting the features, we computed correlated feature sets that had to be perturbed together, as they provide similar, redundant information (**Figure 6A**; see section “Materials and Methods”). Interestingly, for the exact same data, different TFRs provided features with different correlation profiles. The most uncorrelated feature space is generated by the ChoiW representation, while the SLT generates a feature space whose correlation profile is in between the STFT and the CWT.

Figure 6B displays the increase in MSE and magnitude of the decrease in prediction accuracy for the different TFRs on the EEG data. In all representations, the time-frequency components that contribute most to distinguishing between “seen” and “nothing” trials are localized around the alpha and low beta frequency bands (8–15 Hz). The clearest picture is provided by the SLT, which also had the highest classification performance on this data. Interestingly, in addition to the alpha band, the prediction accuracy decrease also indicates that a 47 Hz gamma burst, located around ~ 400 ms was found useful for classification. For the other representations, the distinguishing TFR components were either less localized in frequency, or scattered more toward the sides of the analyzed time segment. In some cases, the analysis also revealed confounding distractors (negative values in **Figure 6B**), especially for the CWT and ChoiW. Taken together,



these results indicate that the most prominent oscillatory features distinguishing “seen” from “nothing” trials are oscillatory bursts in the alpha / lower-beta bands. This is most clearly evidenced by the SLT representation, which also seems to be the most informative (see **Figure 5**).

DISCUSSION

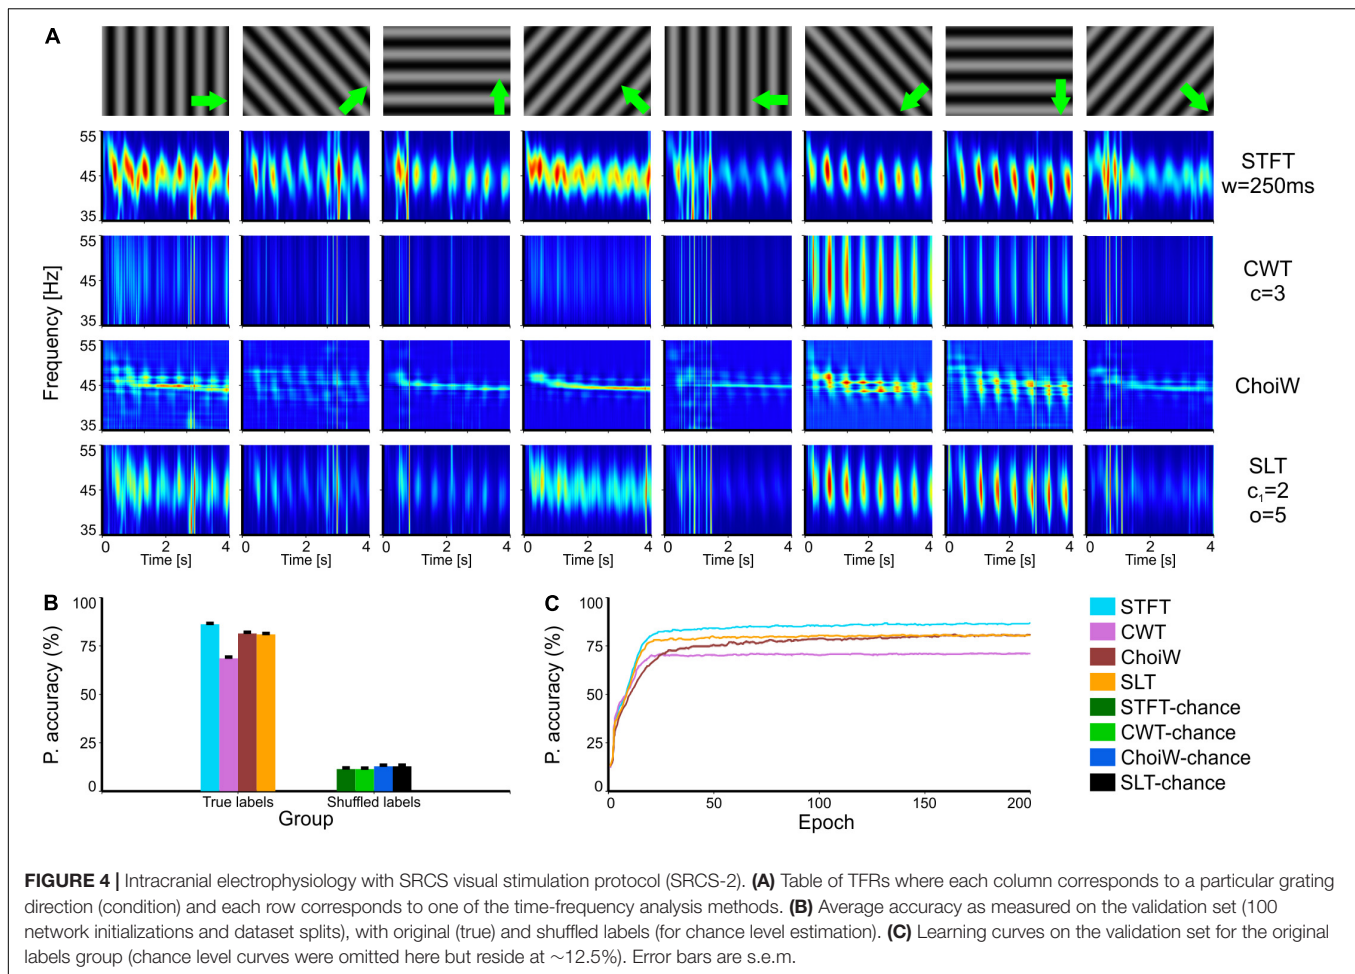
Traditional methods for evaluating the quality of a TFR usually rely on the concept of resolution (Cohen, 1995). In a classical sense, resolution is calculated by estimating the time and frequency uncertainty product of known signals (sine waves, Gaussian atoms). The definition of TFR resolution has been, however, recently challenged (Flandrin et al., 1994; Folland and Sitaram, 1997; Stanković, 2001; Boashash and Sucic, 2003), as this needs to reflect not only the ability to localize an isolated oscillation packet but also the degree to which multiple oscillations can be resolved simultaneously, when they are present in a signal in close temporal and frequency proximity (Moca et al., 2021).

In feature-rich neuroscience data, resolution is not useful for the quantification of the quality of a TFR because the real time-frequency characteristics of the signals cannot be known.

The estimations that we use are only reliable insofar as they unravel as many individual components in the signal as possible, while having as little faults (leakage, cross-terms) as possible.

Here, we have shown that, rather than relying on subjective evaluations of how “sharp” a TFR representation is, one can use automatic methods, such as machine learning, which can discover the relation between the features of brain activity, revealed by different representations, and the particular experimental condition. Although time and frequency resolution may be determining factors for the prediction accuracy and learning efficacy, resolution is clearly not the only relevant factor when judging the quality of a TFR. In particular, it is interesting that techniques which are known to provide the best theoretical time-frequency resolution, such as those derived from the WVD, sometimes underperform in terms of information content on neuroscience data. In that sense, machine learning comes to confirm, in an objective manner, previous empirical results (Moca et al., 2021).

Machine learning techniques for the exploration of brain signals have witnessed increased adoption in recent years (Moca et al., 2009; Vu et al., 2018; Glaser et al., 2019; Roy et al., 2019; Li et al., 2020; Lo Giudice et al., 2022). However, it is important to note that such tools have their own pitfalls and



may present traps that can mislead even the experimented user. For example, the accuracy reached by a trained neural network model does not depend only on the properties of the data, but also on the particular architecture used, training algorithm, the initial weights, or the random splitting into train and test sets (Ciuparu et al., 2020). Therefore, in general, no conclusion can be drawn on a single trained model but multiple models need to be evaluated. In general, the architecture (often chosen empirically) and training algorithm are fixed for a particular problem, but the impact of the random initial conditions can be averaged out by performing multiple, hundreds or thousands of model instantiations, as we have shown here.

Such multi-model approaches are paramount for datasets encountered in neuroscience, which typically have more features than samples (Ciuparu et al., 2020) and exhibit sample numbers usually in the range of tens to hundreds, at best. Training models on such data is almost always cumbersome. For robust training of machine learning models, especially when using deep learning, hundreds of thousands to millions of samples (Deng et al., 2009) are ideal. In neuroscience, the number of available samples is many orders of magnitude below that. While recording more data is rarely possible, one can use techniques involving multiple random splitting and random network initializations, as well as

specialized activation functions that fare well under such difficult conditions (Ciuparu et al., 2020).

Here we have used exclusively MLPs for machine learning, but other options may also be explored. Since TFRs are two-dimensional data representations, it may be interesting to study how convolutional neural networks (CNNs) (LeCun et al., 2015) would fare on such data. There are, however, at least two important aspects to consider. First, CNNs typically require very large datasets to converge, such that convolution kernels can represent the meaningful spatial statistical properties of the data (Krizhevsky et al., 2012). For neural data, it is rarely possible to obtain datasets with even thousands of trials, let alone large datasets with hundreds of thousands / millions of trials. Data augmentation techniques, combined with pooling of datasets across large historical databases may represent a solution. Second, natural images have specific statistical properties (Torralla and Oliva, 2003) that enable CNNs to develop convolution kernels which detect oriented contours, blobs, etc. It is unclear if TFRs possess such universal statistical properties—more likely, TFRs of different types of data may display very different statistics, leading to convolution kernels with very different distributions. In addition, the TFR estimation method is also likely to leave its fingerprint on the statistical properties of the TFR, thereby

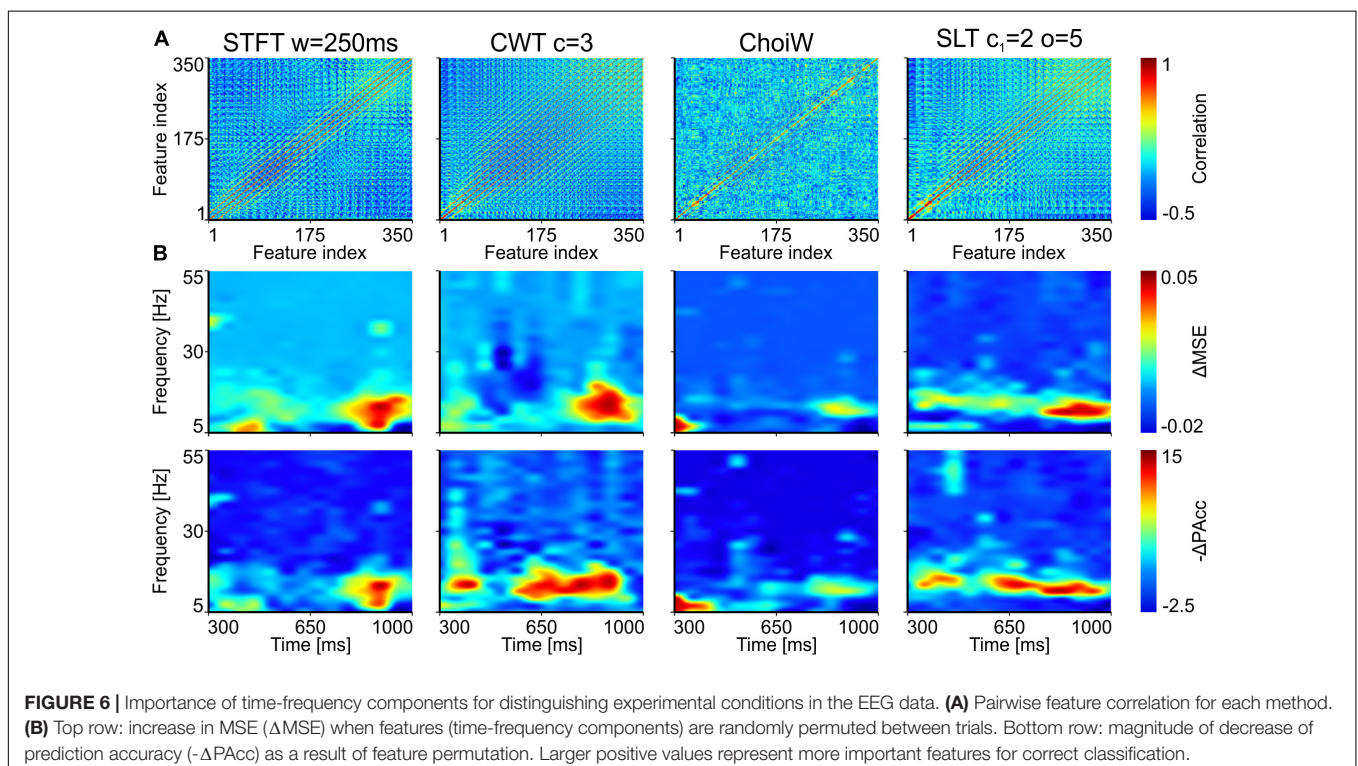
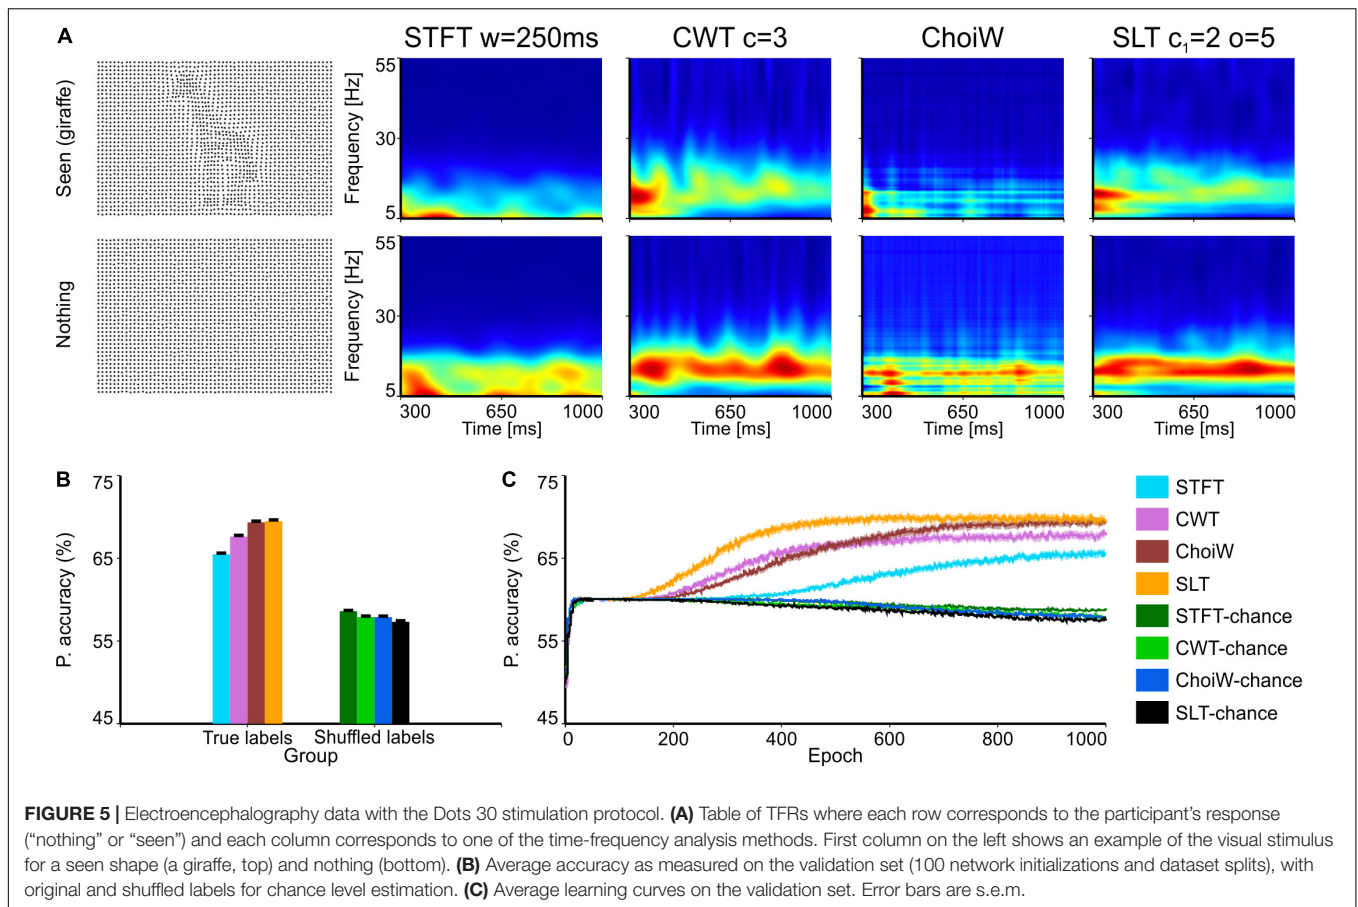


TABLE 1 | Architecture of MLP networks, properties of datasets, and parameters of the training algorithm.

	Parameter name / Property	Value
MLP achitecture for LFP data	Activation function	Soft++, $c = 30$, $k = 1$ (Ciuparu et al., 2020)
	Input layer size	RF: 600, SRCS-1 and 2: 800
	Number of hidden layers	2
	Hidden layer sizes	20% and 10% of input layer size
	Output layer size	8
	Softmax output layer	Yes
	Weight initialization	LeCun normal (LeCun et al., 2012)
	Initial weight range	$[-0.1, +0.1]$
MLP architecture for EEG data	Activation function	Soft++, $c = 30$, $k = 1$
	Input layer size	350
	Number of hidden layers	2
	Hidden layer sizes	20% and 10% of input layer size
	Output layer size	2
	Softmax output layer	Yes
	Weight initialization	LeCun normal
	Initial weight range	$[-0.01, +0.01]$
LFP datasets	Sample dimensionality (size of feature vectors)	RF: 600 (30 time \times 20 freq.), SRCS-1 and 2: 800 (40 time \times 20 freq.)
	Number of samples	80
	Number of classes	8 (directions covering 360° in steps of 45°)
	Training set size	48
	Validation/Test set size	32
	Sample (feature vector) normalization	Min-Max range globally scaled to $[-3.3823, +1.323]$ across entire dataset
EEG dataset	Sample dimensionality	350 (14 time \times 25 freq.)
	Number of samples	150
	Number of classes	2 ("seen" / "nothing")
	Training set size	90
	Validation/Test set size	60
	Sample normalization	Z-scoring each sample
Training parameters	Training algorithm	Batch, IRPROP+ (Igel and Hüsken, 2000)
	Δ_0	0.001 / <i>Fan-in</i>
	Δ_{min}	10^{-15}
	Δ_{max}	0.01 for LFP data, 0.001 for EEG data
	η^-	0.95
	η^+	1.05
	Dropout probability	10% (Hinton et al., 2012)

biasing the learning of the CNN. Nevertheless, exploring the use of CNNs for the evaluation of information content in TFRs is an interesting avenue, worth pursuing in future studies.

The fact that machine learning is able to discover information in neural data is not sufficient for most scientific endeavors. It is equally important to discover how such information is reflected in the properties of the neural data, giving insights into mechanistic brain processes that support perception, cognition, or behavior. As we have shown here, explainable AI techniques (Došilović et al., 2018; Barredo Arrieta et al., 2020; Linardatos et al., 2020) are able to provide insights into why a machine learning model can learn and to indicate where and how relevant

information is expressed in the neural data. This is valuable because such automatic techniques may reveal phenomena that the experimenter could not imagine *a priori*.

For the case of neural oscillations, investigated here, our results indicate that different TFRs can reveal different relevant aspects of the data. The informative features of the TFRs we have tested here show remarkable similarity but also slight differences. The SLT emerges as one of the most informative representations, with good concentration of power and less sensitivity to noise. On the other hand, WVD-based techniques or the STFT can also provide informative representations but tend to be less appropriate for signals with a more complex spectral landscape,

such as the EEG. In all cases, the explainable AI techniques reveal interesting oscillatory components that largely overlap between representations. However, the same tools also indicate that some TFRs can “catch” informative oscillation packets that the others cannot—and this depends not only on the representation but also on the properties of the data.

To conclude, we have shown that machine learning can be a valuable tool for evaluating the quality of TFRs on various types of neural data. In addition, explainable AI techniques can identify what features of the data are relevant and hint at important neural processes that are difficult to discover without automatic tools. The large-scale adoption of such methods is likely to provide a significant boost to the analysis of neural data in the years to come.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author/s.

ETHICS STATEMENT

The studies involving human participants were reviewed and approved by the Local Ethics Committee (1/CE/08.01.2018) and all experiments were conducted in accordance with Directive (EU) 2016/680 and Romanian Law 190/2018. The

patients/participants provided their written informed consent to participate in this study. The animal study was reviewed and approved by the Local Ethics Committee (3/CE/02.11.2018), the National Sanitary and Veterinary Authority (approval 147/04.12.2018), performed in accordance with Directive 2010/63/EU and Romanian Law 43/2014.

AUTHOR CONTRIBUTIONS

HB, VVM, and RCM designed and implemented the methods. HB, A-MI, and RCM recorded the LFP and EEG data. All authors wrote the manuscript.

FUNDING

The research leading to these results has received funding from: NO Grants 2014–2021, under Project contract number 20/2020 (RO-NO-2019-0504), four grants from the Romanian National Authority for Scientific Research and Innovation, CNCS-UEFISCDI (codes PN-III-P2-2.1-PED-2019-0277, PN-III-P3-3.6-H2020-2020-0109, ERA-Net NEURON UNSCRAMBLY, ERA-Net FLAG-ERA ModelDXConsciousness), and a H2020 grant funded by the European Commission (grant agreement 952096, NEUROTWIN). This work was also supported by the Project “Entrepreneurial competences and excellence research in doctoral and postdoctoral programs—ANTREDOC,” project co-funded by the European Social Fund.

REFERENCES

- Aladjalova, N. A. (1957). Infra-slow rhythmic oscillations of the steady potential of the cerebral cortex. *Nature* 179, 957–959. doi: 10.1038/179957a0
- Arnaut, L. G., and Ibáñez, S. (2020). Self-sustained oscillations and global climate changes. *Sci. Rep.* 10:11200. doi: 10.1038/s41598-020-68052-9
- Barraza, P., Pérez, A., and Rodríguez, E. (2020). Brain-to-brain coupling in the gamma-band as a marker of shared intentionality. *Front. Hum. Neurosci.* 14:295. doi: 10.3389/fnhum.2020.00295
- Barredo Arrieta, A., Diaz-Rodríguez, N., Del Ser, J., Bannetot, A., Tabik, S., Barbado, A., et al. (2020). Explainable Artificial intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* 58, 82–115. doi: 10.1016/j.inffus.2019.12.012
- Bârzan, H., Moca, V. V., Ichim, A.-M., and Muresan, R. C. (2021). “Fractional Superlets,” in *Proceedings of the 2020 28th European Signal Processing Conference (EUSIPCO)*. Presented at the 2020 28th European Signal Processing Conference (EUSIPCO), (Amsterdam: EUSIPCO). 2220–2224. doi: 10.23919/Eusipco47968.2020.9287873
- Başar, E. (2013). A review of gamma oscillations in healthy subjects and in cognitive impairment. *Int. J. Psychophysiol.* 90, 99–117. doi: 10.1016/j.ijpsycho.2013.07.005
- Başar, E., Başar-Eroglu, C., Karakaş, S., and Schürmann, M. (2001). Gamma, alpha, delta, and theta oscillations govern cognitive processes. *Int. J. Psychophysiol.* 39, 241–248. doi: 10.1016/S0167-8760(00)00145-8
- Belle, M. D. C., and Diekmann, C. O. (2018). Neuronal oscillations on an ultra-slow timescale: daily rhythms in electrical activity and gene expression in the mammalian master circadian clockwork. *Eur. J. Neurosci.* 48, 2696–2717. doi: 10.1111/ejn.13856
- Boashash, B. (2015). *Time-Frequency Signal Analysis and Processing: A Comprehensive Reference*, 2 Edn. Cambridge, MA: Academic Press.
- Boashash, B., and Susic, V. (2003). Resolution measure criteria for the objective assessment of the performance of quadratic time-frequency distributions. *IEEE Trans. Signal Process.* 51, 1253–1263. doi: 10.1109/TSP.2003.810300
- Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford: Oxford University Press.
- Buzsáki, G. (2002). Theta oscillations in the hippocampus. *Neuron* 33, 325–340. doi: 10.1016/S0896-6273(02)00586-X
- Choi, H.-I., and Williams, W. J. (1989). Improved time-frequency representation of multicomponent signals using exponential kernels. *IEEE Trans. Acoust. Speech Signal Process.* 37, 862–871. doi: 10.1109/ASSP.1989.28057
- Chrobok, L., Belle, M. D. C., and Myung, J. (2021). From fast oscillations to circadian rhythms: coupling at multiscale frequency bands in the rodent subcortical visual system. *Front. Physiol.* 12:738229. doi: 10.3389/fphys.2021.738229
- Ciuparu, A., and Mureşan, R. C. (2016). Sources of bias in single-trial normalization procedures. *Eur. J. Neurosci.* 43, 861–869. doi: 10.1111/ejn.13179
- Ciuparu, A., Nagy-Dăbâcan, A., and Mureşan, R. C. (2020). Soft++, a multi-parametric non-saturating non-linearity that improves convergence in deep neural architectures. *Neurocomputing* 384, 376–388. doi: 10.1016/j.neucom.2019.12.014
- Cohen, L. (1995). *Time-Frequency Analysis: Theory and Applications*. New Jersey, NJ: Prentice-Hall.
- Cordero, E., de Gosson, M., and Nicola, F. (2018). On the reduction of the interferences in the born-jordan distribution. *Appl. Comput. Harmon. Anal.* 44, 230–245. doi: 10.1016/j.acha.2016.04.007
- Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999). Fast network oscillations in the hippocampal ca1 region of the behaving rat. *J. Neurosci.* 19, RC20–RC20. doi: 10.1523/JNEUROSCI.19-16-j0001.1999
- Debnath, L. (2002). “The Wigner-Ville Distribution and Time-Frequency Signal Analysis,” in *Wavelet Transforms and Their Applications*, ed. L. Debnath (Basel: Birkhäuser Verlag). 307–360. doi: 10.1007/978-1-4612-0097-0_5

- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). "ImageNet: a large-scale hierarchical image database," in *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition. Presented at the 2009 IEEE Conference on Computer Vision and Pattern Recognition*, (New Jersey, NJ: IEEE), 248–255. doi: 10.1109/CVPR.2009.5206848
- Došilović, F. K., Brčić, M., and Hlupić, N. (2018). "Explainable artificial intelligence: a survey," in *Proceedings of the 2018 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO). Presented at the 2018 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, (Taiwan: MIPRO), 0210–0215. doi: 10.23919/MIPRO.2018.8400040
- Einevoll, G. T., Kayser, C., Logothetis, N. K., and Panzeri, S. (2013). Modelling and analysis of local field potentials for studying the function of cortical circuits. *Nat. Rev. Neurosci.* 14, 770–785. doi: 10.1038/nrn3599
- Fisher, A., Rudin, C., and Dominici, F. (2019). All models are wrong, but many are useful: learning a variable's importance by studying an entire class of prediction models simultaneously. *J. Mach. Learn. Res.* 20, 1–81.
- Flandrin, P., Baraniuk, R. G., and Michel, O. (1994). "Time-frequency complexity and information," in *Presented at the Proceedings of ICASSP '94. IEEE International Conference on Acoustics, Speech and Signal Processing*, 3, 329–332. Adelaide, SA. doi: 10.1109/ICASSP.1994.390031
- Foffani, G., Priori, A., Egidio, M., Rampini, P., Tamma, F., Caputo, E., et al. (2003). 300-Hz subthalamic oscillations in Parkinson's disease. *Brain* 126, 2153–2163. doi: 10.1093/brain/awg229
- Folland, G. B., and Sitaram, A. (1997). The uncertainty principle: a mathematical survey. *J. Fourier Anal. Appl.* 3, 207–238. doi: 10.1007/BF02649110
- Gabor, D. (1946). Theory of communication. part 1: the analysis of information. *J. Inst. Electr. Eng.* 93, 429–441. doi: 10.1049/ji-3-2.1946.0074
- Glaser, J. I., Benjamin, A. S., Farhoodi, R., and Kording, K. P. (2019). The roles of supervised machine learning in systems neuroscience. *Prog. Neurobiol.* 175, 126–137. doi: 10.1016/j.pneurobio.2019.01.008
- Gourévitch, B., Martin, C., Postal, O., and Eggermont, J. J. (2020). Oscillations in the auditory system and their possible role. *Neurosci. Biobehav. Rev.* 113, 507–528. doi: 10.1016/j.neubiorev.2020.03.030
- Heisenberg, W. (1985). "Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik" in *Original Scientific Papers Wissenschaftliche Originalarbeiten, Werner Heisenberg Gesammelte Werke Collected Works*, eds W. Blum, H. Rechenberg, and H.-P. Dürr (Berlin: Springer), 478–504. doi: 10.1007/978-3-642-61659-4_30
- Herrmann, C. S., and Demiralp, T. (2005). Human EEG gamma oscillations in neuropsychiatric disorders. *Clin. Neurophysiol.* 116, 2719–2733. doi: 10.1016/j.clinph.2005.07.007
- Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. R. (2012). Improving neural networks by preventing co-adaptation of feature detectors. [Preprint] *arXiv:1207.0580 [cs]* doi: 10.1109/TNNLS.2017.2750679
- Hughes, J. R. (2008). Gamma, fast, and ultrafast waves of the brain: their relationships with epilepsy and behavior. *Epilepsy Behav.* 13, 25–31. doi: 10.1016/j.yebeh.2008.01.011
- Hughes, S. W., Lőrincz, M. L., Parri, H. R., and Crunelli, V. (2011). "Chapter 10 - Infraslow (<0.1Hz) oscillations in thalamic relay nuclei: basic mechanisms and significance to health and disease states," in *Progress in Brain Research, Slow Brain Oscillations of Sleep, Resting State and Vigilance*, eds E. J. W. Van Someren, Y. D. Van Der Werf, P. R. Roelfsema, H. D. Mansvelder, and F. H. Da Silva (Amsterdam: Elsevier), 145–162. doi: 10.1016/B978-0-444-53839-0.0010-7
- Huisman, J., and Weissing, F. J. (1999). Biodiversity of plankton by species oscillations and chaos. *Nature* 402, 407–410. doi: 10.1038/46540
- Igel, C., and Hüsken, M. (2000). "Improving the Rprop Learning Algorithm," in *Proceedings of the Second International ICSC Symposium on Neural Computation (NC 2000)*, (Cambridge, MA: ICSC Academic Press), 115–121.
- Keil, A., Müller, M. M., Ray, W. J., Gruber, T., and Elbert, T. (1999). Human gamma band activity and perception of a gestalt. *J. Neurosci.* 19, 7152–7161. doi: 10.1523/JNEUROSCI.19-16-07152.1999
- Klimesch, W. (1999). EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain Res. Rev.* 29, 169–195. doi: 10.1016/S0165-0173(98)00056-3
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). "ImageNet classification with deep convolutional neural networks," in *Proceedings of the 25th International Conference on Neural Information Processing Systems (NIPS'12)*, eds F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Lake Tahoe, NV: Curran Associates, Inc), 1097–1105.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521, 436–444. doi: 10.1038/nature14539
- LeCun, Y. A., Bottou, L., Genevieve, B., and Müller, K.-R. (2012). "Efficient backprop," in *Neural Networks: Tricks of the Trade: 2nd Edition, Lecture Notes in Computer Science*, eds G. Montavon, B. Geneviève, and K.-R. Müller (Berlin: Springer), 9–48. doi: 10.1007/978-3-642-35289-8_3
- Lega, B. C., Jacobs, J., and Kahana, M. (2012). Human hippocampal theta oscillations and the formation of episodic memories. *Hippocampus* 22, 748–761. doi: 10.1002/hipo.20937
- Li, G., Lee, C. H., Jung, J. J., Youn, Y. C., and Camacho, D. (2020). Deep learning for EEG data analytics: a survey. *Concurr. Comput.* 32:e5199. doi: 10.1002/cpe.5199
- Linardatos, P., Papastefanopoulos, V., and Kotsiantis, S. (2020). Explainable AI: a review of machine learning interpretability methods. *Entropy* 23:18. doi: 10.3390/e23010018
- Lo Giudice, M., Varone, G., Ieracitano, C., Mammone, N., Tripodi, G. G., Ferlazzo, E., et al. (2022). Permutation entropy-based interpretability of convolutional neural network models for interictal eeg discrimination of subjects with epileptic seizures vs. psychogenic non-epileptic seizures. *Entropy* 24:102. doi: 10.3390/e24010102
- Melloni, L., Molina, C., Pena, M., Torres, D., Singer, W., and Rodriguez, E. (2007). Synchronization of neural activity across cortical areas correlates with conscious perception. *J. Neurosci.* 27, 2858–2865. doi: 10.1523/JNEUROSCI.4623-06.2007
- Moca, V. V., Bârzan, H., Nagy-Dăbăcan, A., and Mureșan, R. C. (2021). Time-frequency super-resolution with superlets. *Nat. Commun.* 12:337. doi: 10.1038/s41467-020-20539-9
- Moca, V. V., Scheller, B., Mureșan, R. C., Daunderser, M., and Pipa, G. (2009). EEG under anesthesia—feature extraction with TESPARE. *Comput. Methods Programs Biomed.* 95, 191–202. doi: 10.1016/j.cmpb.2009.03.001
- Moca, V. V., Țincaș, I., Melloni, L., and Mureșan, R. C. (2011). Visual exploration and object recognition by lattice deformation. *PLoS ONE* 6:e22831. doi: 10.1371/journal.pone.0022831
- Mureșan, R. C., Jurjuț, O. F., Moca, V. V., Singer, W., and Nikolić, D. (2008). The oscillation score: an efficient method for estimating oscillation strength in neuronal activity. *J. Neurophysiol.* 99, 1333–1353. doi: 10.1152/jn.00772.2007
- Nikolić, D., Mureșan, R. C., Feng, W., and Singer, W. (2012). Scaled correlation analysis: a better way to compute a cross-correlogram. *Eur. J. Neurosci.* 35, 742–762. doi: 10.1111/j.1460-9568.2011.07987.x
- Oostenveld, R., Fries, P., Maris, E., and Schoffelen, J.-M. (2010). FieldTrip: open source software for advanced analysis of meg, eeg, and invasive electrophysiological data. *Comput. Intell. Neurosci.* 2011:e156869. doi: 10.1155/2011/156869
- Roy, Y., Banville, H., Albuquerque, I., Gramfort, A., Falk, T. H., and Faubert, J. (2019). Deep learning-based electroencephalography analysis: a systematic review. *J. Neural Eng.* 16:051001. doi: 10.1088/1741-2552/ab260c
- Skaggs, W. E., McNaughton, B. L., Wilson, M. A., and Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6, 149–172. doi: 10.1002/(SICI)1098-1063(1996)6:2<149::AID-HIPO6<3.0.CO;2-K
- Stanković, L. (2001). A measure of some time–frequency distributions concentration. *Signal Proc.* 81, 621–631. doi: 10.1016/S0165-1684(00)00236-X
- Steriade, M. (2003). The corticothalamic system in sleep. *Front. Biosci.* 8, 878–899. doi: 10.2741/1043
- Steriade, M., and Amzica, F. (1998). Slow sleep oscillation, rhythmic K-complexes, and their paroxysmal developments. *J. Sleep Res.* 7, 30–35. doi: 10.1046/j.1365-2869.7.s1.4.x
- Stuiver, M., and Braziliunas, T. F. (1989). Atmospheric ¹⁴C and century-scale solar oscillations. *Nature* 338, 405–408. doi: 10.1038/338405a0
- Tal, I., Neymotin, S., Bickel, S., Lakatos, P., and Schroeder, C. E. (2020). Oscillatory bursting as a mechanism for temporal coupling and information coding. *Front. Comput. Neurosci.* 14:82. doi: 10.3389/fncom.2020.00082

- Teleńczuk, B., Dehghani, N., Quyen, M., Cash, S. S., Halgren, E., Hatsopoulos, N. G., et al. (2017). Local field potentials primarily reflect inhibitory neuron activity in human and monkey cortex. *Sci. Rep.* 7:40211. doi: 10.1038/srep40211
- Torralba, A., and Oliva, A. (2003). Statistics of natural image categories. *Network* 14, 391–412. doi: 10.1088/0954-898x_14_3_302
- Traub, R. D., Hawkins, K., Adams, N. E., Hall, S. P., Simon, A., and Whittington, M. A. (2020). Layer 4 pyramidal neuron dendritic bursting underlies a post-stimulus visual cortical alpha rhythm. *Commun. Biol.* 3, 1–13. doi: 10.1038/s42003-020-0947-8
- Vu, M.-A. T., Adalı, T., Ba, D., Buzsáki, G., Carlson, D., Heller, K., et al. (2018). A shared vision for machine learning in neuroscience. *J. Neurosci.* 38, 1601–1607. doi: 10.1523/JNEUROSCI.0508-17.2018
- Wilson, M. A., and McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science* 265, 676–679. doi: 10.1126/science.8036517

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Bârzan, Ichim, Moca and Mureşan. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.