MINISTRY OF EDUCATION

TECHNICAL UNIVERSITY
OF CLUJ-NAPOCA, ROMANIA

Electronics, Telecommunications and Informational Technologies Engineering

# PhD THESIS
## - ABSTRACT -

# Fast oscillations in cortical circuits - from estimation to therapeutical and practical applications

**PhD Student:**
**Harald Bârzan**

**PhD Supervisor:**
**Prof. Eng. Corneliu RUSU, PhD**

**Examination committee:**
Chair:
Prof. Eng. **Sorin Hintea**, PhD - Technical University of Cluj-Napoca
PhD supervisor:
Prof. Eng. **Corneliu Rusu**, PhD - Technical University of Cluj-Napoca
Members:
- Prof. Eng. **Dinu Coltuc**, PhD - "Valahia" University of Târgoviște
- Prof. Eng. **Adrian Graur**, PhD - "Ștefan cel Mare" University of Suceava
- Prof. Eng. **Dorin Petreuș**, PhD - Technical University of Cluj-Napoca

**- Cluj-Napoca -**
**2022**

str. Memorandumului nr. 28, 400114 Cluj-Napoca, România
tel. +40-264-401200, fax +40-264-592055, secretariat tel. +40-264-202209, fax +40-264-202280
www.utcluj.ro

# 1 Introduction

Across all of existence there is probably no self-contained system, biological or otherwise, that can rival the mammalian brain in complexity. Of all species on Earth, present or extinct, we humans have made use of it the most. From crafting tools to build, hunt and forage, to inventing language, society and then civilization, we became the dominant species on the planet mostly due to this single organ. It would stand to reason that we would know a lot about it, but alas this is not the case - extensive brain research only started in the mid 20th century.

A lot of work went into neuroscience in the last few decades. Here is a brief summary of what we have uncovered thus far: we looked at brains under the microscope and found out that they're built out of two types of cells, neurons and glia, with the former having an active role in brain function (cognition, memory etc) and the latter having a more supportive role. We measured electrical activity at different levels of organization (from scalp electrodes to single cell recordings) and learned that neurons conduct signals through electrical potentials - waves through the dendrites and impulses through the axons. We also found that neurons synchronize their activity, giving rise to brain*waves*.

We broke down the neurons' biochemistry and found that they communicate with each other using synapses in which messenger molecules activate a target neuron *via* a lock-key interaction. These messenger molecules have varied effects on the neurons they target and these effects mainly depend on the kind of receptors the molecules bind to - they can excite or inhibit a target neuron or they can cause more subtle metabolic changes to occur inside it. The topmost classification of neurons is excitatory and inhibitory neurons, with an approximate 80-20% distribution.

With the invention of more modern recording and imaging techniques, we managed to establish that the mammalian brain has a topological organization, in which different brain areas specialize in different tasks. For example, the hypothalamus controls involuntary bodily functions such as breathing and heartbeat, while the cortex is more preoccupied with more complex and often voluntary behavior: cognition, memory, planning and so on. The cortex itself is thought of as the main reason behind humans' and other primates' intellect.

An interesting topic of research, coincidentally the centerpiece of this thesis, is the domain of neural oscillations. These rhythmic patterns occur in both cortical and subcortical areas and, for particularly fast oscillations, are thought to arise due to interactions between excitatory and inhibitory neurons, in which inhibition is rhythmically dialed up and down, providing excitatory neurons with windows of opportunity in which their firing can affect downstream neurons. Numerous studies have gone into establishing whether the oscillatory behavior of neural circuits is a geniune mechanism governing brain function or if it is just an epiphenomenon, caused by a more intricate mechanism.

The thesis summarized herein focuses on building methods to investigate oscillatory activity in neural time series data, with an emphasis on a new time-frequency analysis method called *Superlet Transform* (SLT), based on the commonly used continuous wavelet transform (CWT). This new technique effortlessly produces high time-frequency resolution scalograms using an information-theoretic approach of *minimum-mean cross-entropy* (geometric averaging) to combine multiple CWT scalograms into a single, high-resolution representation.

The thesis also shows two different SLT applications in scenarios where neural oscillations are concerned. First, a putative therapy for Alzheimer's disease called GENUS (Gamma ENtrainment Using Sensory stimuli) is explored. First published by a team led by Dr. Li-Huei Tsai, this therapeutic approach focuses on inducing gamma oscillations (30-80 Hz) using flickering light and claims that this gamma entrainment helps neurons and surrounding glial cells to break apart and destroy the amyloid-beta plaques that are the hallmark of this disease, while also cutting back on the genetic expression (synthesis) of these plaques. Here, superlets are used to assess

the properties of the gamma entrainment generated by the flickering light in local field potentials (LFPs, continuous time series signals) in mice, acquired by using intracortical multi-channel silicon probes.

The second scenario involves gamma oscillations that are not entrained by an outside stimulus, but produced willingly by the subject. These rhythms are picked up using a non-invasive electroencephalography (EEG) setup and are used as control signals for a brain-computer interface (BCI). Gamma oscillations were typically ignored for BCI applications using EEG, mainly because these oscillations are low amplitude events with a typically low spatial footprint, making them very difficult to observe in EEG recordings where electrode sizes are large.

In what follows, the methods used in this thesis (the recording and analysis software, the superlets) and the results that they produced will be summarized.

# 2 Methods

There are a number of methods discussed throughout the thesis and they should be conceptually split in two categories: methods that were available to the scientific community before this thesis and methods that have been introduced with this thesis.

Filling the ranks of the first category are the usual suspects: the decades-old techniques for both intracortical recordings, using gold-plated silicon probes physically piercing the cortex through a craniotomy, and non-invasive scalp recordings using EEG. Pair these signal acquisition methods with signal processing techniques such as analog filtering, amplification, analog-to-digital conversion and storage. For properly processing and analyzing neural time series data, digital filtering (such as IIR filtering) is needed, as is time-frequency spectral analysis using variations of the Fourier and Wavelet transforms. Cross-correlation analysis and spectral coherence analysis are also common, especially when dealing with local field potentials, where interareal communication through time or frequency correlations can be assessed. All these techniques are part of the typical recording analysis workflow a neuroscientist would normally deal with, using tools such as LabVIEW, MATLAB, Python.

The second category, which is explained in more detail, concerns the tools and methods introduced in the thesis itself. Here we have a varied portfolio of items ranging from visual stimulation devices for GENUS therapy, a data analysis framework developed in C# with a focus on neural time-series data, a novel time-frequency spectral estimation algorithm (superlets) and subsequent variations on it, recording software, stimulation software and novel machine learning-based analytical procedures for spike detection and power spectrum classification.

## 2.1 Recording neural signals

Two methods for recording neural signals are discussed at length in the thesis, covering two out of the three canonical scales of brain dynamics: mesoscale (small to medium sized neural circuits) and macroscale (whole brain areas) activity. Both methods produce neural time-series data as their output, which can be further processed and analyzed using typical signal processing techniques such as digital filtering and Fourier transforms, respectively.

### 2.1.1 Intracortical recordings

The best way of recording mesoscale electrical activity is to cut away a portion of scalp and skull above the area of interest (craiotomy), and insert some kind of recording device in the exposed neural tissue. This recording device can take multiple forms, depending on experimental demands. In this thesis the *acute silicon probes* are used, where acute means that the experiment has a short time frame, usually a few hours, before the probe is removed. The only exceptions

to acute recordings are the cat recordings performed by Dr. Danko Nikolic at the Max Planck Institute for Brain Research, and the macaque recordings performed by Prof. Dr. Wolf Singer's team at the Ernst Strüngmann Institute, where chronic microelectrode arrays were used instead, where chronic means that the subjects "wear" the recording device for a sizable period of time, fixed to the skullbone with cement (the craniotomy is closed).

Generally speaking, most intracortical recording techniques have a decent number of channels (e.g. 32 channels) with a sufficiently high sampling rate (e.g. 32kSamples/sec). The amplifier-digitizer systems they use are quite complex and precise (resolution of at least 16 bits per sample), have stringent grounding and shielding demands, and necessitate more involved synchronization setups in order to enable event signalling with high temporal precision.

### 2.1.2 Electroencephalography

EEG has the distinction of being one of the oldest recording techniques still ubiquitously used today. The main gist of EEG is that large, synchronized neuron populations give rise to brainwaves of sufficient magnitude such that they can be recorded from outside the scalp, without necessitating any invasive procedure to perform, making it feasible in clinically healthy human subjects. The technique itself covers the whole scalp with equidistantly placed silver chloride electrodes in large numbers (typically 64-256 electrodes).

The drawbacks of EEG when compared to intracranial recordings (of any type) stem from the fact that the skull and scalp form barriers between the source (neural tissue) and sensor. These barriers have a very dramatic effect on the quality of the signal, acting as natural low-pass filters in both temporal frequency and spatial frequency. This means that the signal will only reflect low frequency activity with more fidelity and that the signal will appear smeared across multiple electrodes. More novel digital processing techniques, such as average referencing and the Surface Laplacian, can be used to compensate for some of the losses induced by these barriers.

## 2.2 Neuroscience-related software

A number of pieces of software were developed for use in the lab and have been thus included in this thesis. The centerpiece of this batch of software is the TINS library (stylized as `TINS`).

### 2.2.1 The TINS neural data analysis framework

Development of an in-house data management, processing and analysis toolkit begun as a training exercise in implementing various digital signal processing techniques such as digital filtering, short-time Fourier transforms (STFT), discrete and continuous wavelet transforms (WT), spectral coherence, cross-correlation and so on. As things kept being added to the toolkit, it started to be less like a plug-in for a greater development environment and more like its own thing - a framework, in the computer programming sense.

A lot of development went into this budding framework, which is being continuously updated with neuroscience-related analytical procedure as the need arises. The main part of the framework, the core library (`TINS.Core`), contains various classes and routines used for organizing sets of neural time-series data (commonly stored as 32-bit floating point numbers in binary format), parsing descriptor files, loading and saving data, bookkeeping events and organizing trial structures. The core library also contains all of the usual preprocessing components needed for dealing with continuous signals such as digital filters (low-pass, high-pass, band-pass and notch infinite impulse response filters) and resampling techniques (decimation and interpolation). A wide array of analytical techniques are also incorporated: Fourier transforms (1D, 2D), STFTs, coherence and phase-locking analysis, wavelet and superlet analysis. It even has its own implementation of a deep learning algorithm, a multilayer perceptron (MLP), that is fully configurable

and interfaces well with the rest of the framework. Dimensionality reduction algorithms, such as principal component analysis (PCA) are also present.

In the data management perspective, the beating heart of this framework is a set of containers based on matrices and vectors (`Matrix<T>`, where `T` can be any kind of object, including numbers, strings or other matrices), with basic linear algebra procedures being constantly implemented and optimized.

Aside from the core library, which is meant to be cross-platform using the .NET 6 runtime, there is also a Windows-only charting library (`TINS.Visualization`) that can render line graphs, barplots and heatmaps, which makes creating figures and saving them as PNG image files extremely easy.

### 2.2.2 Recording and stimulation software

Building on top of the TINS framework are also some software that are routinely used in the lab for visual stimulation and intracranial recording.

The first of these is the visual stimulation software. It started as an alternative to a Windows XP software that was previously used and had the job of rendering a series of frames to a computer screen at a preset frame rate (60 Hz) and to use an external digital acquisition board (NI-DAQmx 6636) to relay TTL signals to the amplifier-digitizer system that is actually recording the analog data from the animal subject. This way, the triggers - events, such as *trial start*, *stimulation start*, *trial end* and so on - would be perfectly aligned to the recorded data, for experiments in which the precise timing is critical.

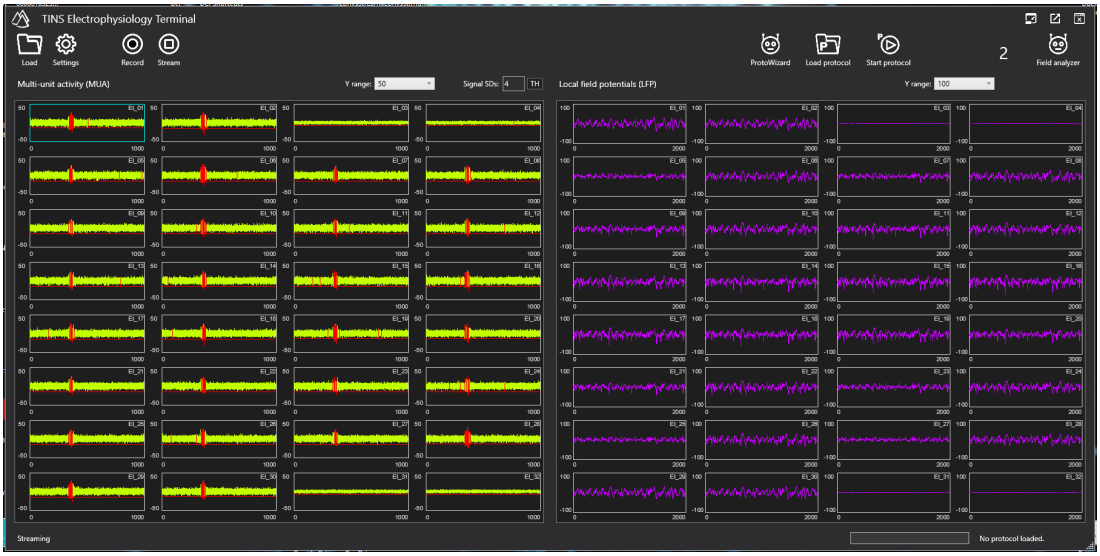


**Figure 1: The TINS Electrophysiology Terminal** software, meant to be used in electrophysiology experiments.

Starting with Windows 7, Microsoft has deprecated the component that was doing the frame rendering for this old software, DirectDraw, in favor of integrating it into its ever more popular DirectX framework. The problem at the time was that for the old software, which was developed in Pascal, there was no easy way to interface with the DirectX libraries, so starting from scratch in another language was warranted. Leveraging the fact that the TINS library already contained the necessary code for reading and writing configuration files and trial information (ETI) files, and the fact that DirectX bindings for C# are common (libraries like Vortice, SharpDX and SlimDX are the most popular), a re-write of this software was possible.

The second, and more extensive, piece of software built to be used in the lab was the *TINS Electrophysiology Terminal*, a fully-sized digital acquisition software to replace an older, phased-out commercial program from Multichannel Systems. The new acquisition software was needed

not only because the old one was deprecated, but also because with it we can fully control all aspects of acquisition *and* stimulation, enabling us to build a closed-loop system in which stimulus can be driven by features of the signals in real time, a feat that was not possible using the older software. Additional perks include the more modern user interface and the fact that this software can directly output the data formats we use in the lab and thus two-step conversions of large datasets are no longer needed.

## 2.3 Superlet Transform

The crux of this thesis revolves around a new time-frequency analysis algorihm, the *superlet transform* (SLT), and its applications in a few neuroscience scenarios. The base equation for the SLT is the following:

$$SLT_{x,c_1,o_f}(f,t) = \left[\prod_{i=1}^{o_f} P_x(c_1 i, f, t)\right]^{\frac{1}{o_f}} \tag{1}$$

where

- $x$ is the input signal $x(t)$,
- $c_1$ is the number of cycles of the shortest wavelet,
- $o_f$ is the order of the superlet at frequency $f$,
- $P_x(c, f, t)$ is the power ($2|A|^2$) of the response $R_x$, i.e. the convolution of the signal with the wavelet:

$$P_x(c,f,t) = 2|R_x(c,f,t)|^2 = 2|(\psi_{f,c} * x)(t)|^2 \tag{2}$$

where, $\psi_{f,c}$ is the Morlet wavelet with central frequency $f$ and number of cycles $c$.

The SLT is an extension of the CWT (see Equation 2), where the power values of the wavelet responses - the results of the convolution operations of the input signal with each wavelet - are geometrically combined to produce a single value per time-frequency point. Each superlet is a set of Morlet wavelets with the same center frequency but increasing number of cycles. The *order* of the superlet defines how many wavelets it contains, while the *base cycles* parameter defines how long the shortest wavelet in the set is.

Briefly, the superlets work by optimizing the resulting response to be as sharp as possible in both time and frequency. The shortest wavelets favor high temporal resolution but poor frequency resolution, while the longest wavelets favor frequency resolution at the cost of temporal resolution. The geometric averaging makes it so that the extremities (shortest and longest wavelets) of the set effectively veto in time and frequency, essentially creating a "bounding box" around the time-frequency point. An *adaptive* SLT (ASLT) was created to make the order of the superlets scale with increase in frequency, sporting even better frequency resolution with increasing frequency, but at the cost of banding - visually observable tearing of the scalogram - at the transitions between the integer valued orders.

A similar method to the SLT, but applied in 2D, is used in super-resolution optical microscopy, in which the response of multiple illumination patterns are geometrically combined to obtain an image resolution that is beyond what can be obtained with classical optical microscopy. It is also important to point out that while the SLT appears to defy the Heisenberg-Gabor uncertainty principle, it is actually not, because the principle applies to *single measurements*, not combined measurements which is what these super-resolution techniques do.

Not having yet a very rigorous mathematical definition for the SLT, we sought to improve it and generalize it. First we sought a *fractional* definition for it, in which the order can be any positive, non-zero real number. This new *fractional SLT*, using weighted geometric averaging

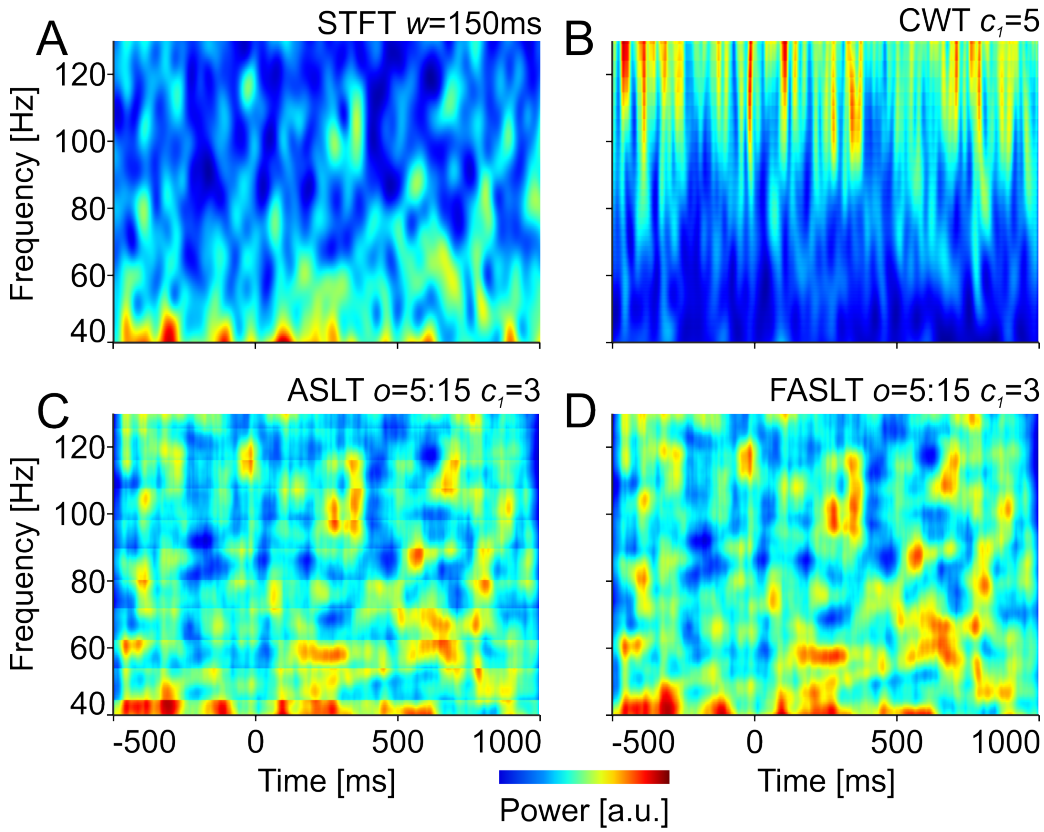


**Figure 2: Examples of ASLT and fractional ASLT against the traditional STFT and CWT** on EEG data with a rich time-frequency landscape. **A.** Mid-sized window STFT showcasing the dilution phenomenon. **B.** CWT with a fixed mid-sized wavelet scale (5 cycles), showcasing the poor frequency resolution at higher frequencies (redundancy). **C.** Adaptive SLT (orders 5 through 15) showing the good time-frequency resolution and absence of dilution through the whole frequency range, but also exhibiting the banding effect. **D.** Fractional adaptive SLT with the same order range as in C, showcasing the elimination of the banding effect. Taken from Bârzan *et al.* 2021.

to include an extra wavelet with subunit weight, solves the banding phenomenon which was inherent to the integer-order ASLT. The equation of the fractional SLT then becomes:

$$FSLT_{x,c_1,o_f}(f,t) = \left[ P_x(c_1(o_i+1),f,t)^{\epsilon} \prod_{i=1}^{o_i} P_x(c_1 i,f,t) \right]^{\frac{1}{o_f}} \tag{3}$$

where, $o_i$ is the integer part of the real order $o_f$ (i.e. $o_i = \text{floor}(o_f)$) and $\epsilon$ is the fractional part (i.e. $o_f - o_i$). $\epsilon$ is the weight of the extra wavelet.

Finally, applying the geometric mean on the polar notations of the complex responses, we created a general equation for the SLT that works for complex inputs and outputs. Earlier, the SLT was only defined for real numbers - a magnitude-squared operation converted each wavelet response into spectral power before geometrically averaging. Thus we now have a general equation for the SLT, which we will display for illustrative purposes:

$$\begin{aligned} SLT_{x,c_1,o_f}(f,t) = &\left[ \mu_x(c_1(o_i+1),f,t)^{\epsilon} \prod_{i=1}^{o_i} \mu_x(c_1 i,f,t) \right]^{\frac{1}{o_f}} \\ &\cdot \exp\left[ j \arg\left( \frac{\epsilon\theta_x(c_1(o_i+1),f,t) + \sum_{i=1}^{o_i} \theta_x(c_1 i,f,t)}{o_f} \right) \right] \end{aligned} \tag{4}$$

where, $\mu_x$ is the magnitude of the wavelet response (i.e. $|R_x|$ from Equation 2) and $\theta_x$ is the *versor* (phase vector of magnitude 1) of the wavelet response, given by $R_x {\mu_x}^{-1}$.

Lastly, we are seeking a rigorous mathematical definition for assessing the phase synchronization of two signals in this superlet framework, similarly to what is done when computing magnitude-squared coherence (MSC). With superlets, we gain the advantage that the analysis window (i.e. the wavelets) is adapted to the frequency of the process being measured, as opposed to using a single size for all frequencies, and the order and base cycles parameters are more adequate because they refer to scale-free processes which are common in neural data. The work described in this thesis referring to *superlet coherence*, the current working name for the new phase synchronization estimation method, is only in its preliminary stage.

## 2.4 ML classification of time-frequency spectra

To compare the newly developed superlets against the more traditional STFT, CWT and reduced interference distributions (RIDs), such as the Choi-Williams (CW) distribution, we needed a way to make the test as realistic as possible, using a few scenarios on real neural data. Multiple data sources such as intracortical mouse recordings and human EEG recordings were used, having multiple types of visual stimuli - trial conditions. An MLP was then trained to classify the resulting time-frequency representations according to the trial conditions. Both maximum accuracy and learning curves were measured.

Additionally, we devised a new feature perturbation method, called *joint feature permutation*, that can determine what features are most relevant for classification. This is done using a special feature permutation scheme: the correlation matrix of the individual features is first computed. Then, for each feature, a correlation threshold is set so that all other correlated features are perturbed together with the feature. This is done in order to avoid the effect of perturbing single features that are strongly correlated with other features, which would show no effect on the predictive accuracy measurement of the classifier. The reason we do this is to avoid using dimensionality reduction, which would eliminate the correspondence between the remaining features and their position on the time-frequency representation. In the end, this allows us to build heatmaps that show us where the informational ”hotspots” are that contribute the most to predictive accuracy and minimize mean squared error.

## 2.5 Spike detection with deep learning

Another aspect of electrophysiological neural data that often warrants a lot of analysis are the neuronal action potentials, informally called *spikes*, electrical impulses that travel from the neuron’s body along the axon to the synapses, signalling the release of neurotransmitters. Spiking activity is very informative in neuroscience, as the temporal widths of the spikes can inform as to what kind of neuron is firing it - shorter spikes are usually from an inhibitory interneuron, while larger spikes usually signal the presence of an excitatory piramidal cell.

The typical workflow in spike extraction is band-pass filtering the signal to eliminate low frequency fluctuations (<300 Hz) and high frequency noise (usually >5000 Hz). This essentially provides us with what is called *multiunit activity* (called this way because there are multiple neurons - units - contributing spikes to the signal), a normally distributed signal with occasional undershoots. These undershoots are the spikes. A detection algorithm is then run to obtain the timestamps and waveforms of the individual spikes. Lastly, the spikes are *sorted* by their characteristics, usually *via* a clustering algorithm, into single units.

This thesis introduces an algorithm that focuses exactly on the detection part. Usually, researchers employ a thresholding algorithm with a fixed threshold set at a multiple of the standard deviation (SD) or interquartile interval (IQI) of the multiunit signal. This is very quick for computers to process - it can be done online while recording (the acquisition software in Figure 1 uses this method) - but it also means that the amplitude distribution of the obtained spikes

will be cut at the threshold, with the implication that there are a lot of spikes that are close to the threshold but are missed by this method.

In order to build a greedier spike detection algorithm, we started with the thresholding method. An MLP classifier with two output nodes (*spike* and *non-spike*) is trained on a dataset formed of the threshold-detected the spikes and an equal number of random bits of signal with no spikes. The trained classifier is then ”slid” across the multiunit signal sample by sample, with the *spike* class output probability at each point saved in a *probability signal*. The probability signal is then thresholded and waveforms are cut from the corresponding region in the multiunit signal.

## 2.6 Brain-computer interface

Using knowledge gained from working with gamma oscillations, EEG and high-resolution time-frequency analysis, we set out to develop an EEG-based brain-computer interface (BCI). Several concepts were borrowed from earlier implementations of BCIs, such as motor imagery, which implies the conceptualization of moving one’s limbs, creating the same patterns of activity in the motor cortex as would the actual act of moving them.

Many kinds of motor imagery experiments exist in the literature using Mu rhythm desynchronization, that is, a drop in spectral power in the alpha band in the motor cortex. However, to estimate power in the alpha band (8-13 Hz) precisely, several cycles of alpha oscillations would have to be measured, meaning a temporal integration window that would be hundreds of milliseconds long, thereby limiting the effective polling rate (i.e. number of commands that can be issued per second) of the BCI. Therefore it was decided to go with a higher frequency range, therefore shortening this temporal integration window.

There are numerous challenges in extracting high-frequency activity (>30 Hz) from EEG signals, as the scalp and skull act as low-pass filters in both space and frequency, meaning that the localization of the high-frequency source is diminished and its magnitude dampened. High-pass spatial filters, such as the Surface Laplacian (SL) are great tools to mitigate these losses, as in brain dynamics oscillation frequency and spatial footprint are inversely correlated - the higher the frequency, the smaller the source neuronal population. SL deals with this by performing a distance-weighted average of the signals of the neighbors of an electrode and subtracting that average activity from the target electrode, thereby acting as a high-pass filter in the spatial domain. Due to the inverse relationship between frequency and space, this also clears out the low frequency activity in the electrodes. A number of other processing steps are performed to improve the quality of the signal.

Superlets are then used to compute the time-frequency representations of the left and right target electrodes (in real time), which would be situated directly above the hand area of the motor cortex. Based on the activity of the left and right electrodes in the gamma range a decision on whether to emit a command and the nature of that command is taken by the BCI.

# 3 Results

This thesis has produced a number of interesting results across all of its explored avenues - signal processing, GENUS therapy and BCI. These results have been published in various scientific journals and conferences.

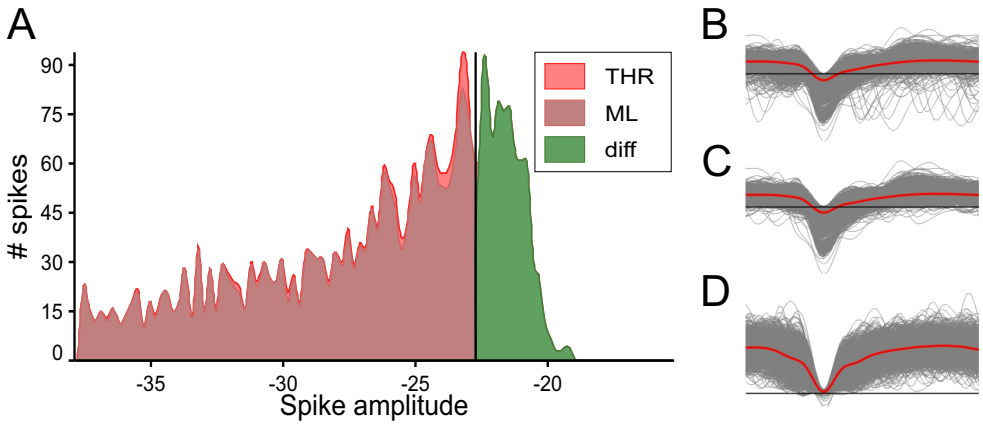


**Figure 3: Spike amplitude distribution and spike waveform averages. A.** Contrast between threshold-based (THR, bright red) and ML-assisted detection (ML, dark red). The green area represents spikes found only by the ML method. Average (red) and individual (gray) shapes of the threshold-detected spikes (**B.**), ML-detected spikes (**C.**) and the set difference between the former two (**D.**). Taken from Bârzan *et al.* 2020.

## 3.1 ML-facilitated spike detection

The ML approach to spike detection has produced very interesting results. Not only does it detect >95% of the spikes detected by the traditional thresholding method (the waveforms of the undetected 5% hardly look like spikes at all), but it also has a novelty rate - percent of extra spikes discovered with respect to the thresholding method - of over 15%, meaning that at least 15% of all spikes are missed by the thresholding method.

Lowering the detection threshold helps up to a point at which the thresholds ”dig into” the *hash* - the richer part of the multiunit signal’s distribution, resulting in more spikes of lower quality. The ML method overcomes this by looking at the shape of the impulse as well as its amplitude. The result is a more natural-looking spike amplitude distribution that tapers off with increasing amplitude, instead of featuring a brick wall-like cutoff (Figure 3). This is also confirmed by the natural shape of the spikes in Figure 3D, showing the novel spikes exhibiting the canonical spike features.

## 3.2 Superlet Transform

The SLT has captivated a lot of people since its introduction judging by citation metrics and GitHub activity. At the time of writing, it has produced 3 scientific papers from the authors, papers that have drawn attention from the scientific and engineering communities. In what follows a number of interesting findings will be highlighted.

### 3.2.1 Superlets allow for robust time-frequency super-resolution

The first striking aspects when comparing multiple time-frequency representations obtained with different methods (STFT spectrograms, CWT/SLT scalograms, reduced interference distributions and so on), is how concentrated the representations for the SLT are.

In one example, we generate a surrogate signal where 3 sets of sine wave with time- and frequency neighbors and use it to compute the representations. What we observe in these representations is that the CWT has drastically poor frequency resolution and really good time resolution, while the STFT suffers less from this problem but has worse temporal resolution. The

SLT, with its geometric combination of wavelets optimized for different resolutions, gets rid of this problem. Here, both time and frequency resolution are good. Figure 3 shows a comparison between STFT, CWT and (F)ASLT, showcasing the latter's improved time-frequency resolution.

Another interesting aspect is how the SLT is able to extract oscillation packets buried deep in brain noise. In one experiment, we added a few oscillatory packets to one trial out of 84 EEG trials. STFT, MMCE, CWT and SLT were used to attempt to identify these packets and, out of all these methods, the SLT was the most effective in extracting them.

### 3.2.2 Fractional superlets solve ASLT's main drawback

Fractional SLT was introduced in an early effort to generalize the formula for SLT for positive real orders. It also helps by eliminating one of ASLT's most important drawbacks from the original Nature Communications paper: banding. This phenomenon appeared in ASLT due to the fact that the order was defined as integers, and transitions between integer orders causes a tearing-like effect visible throughout the frequency range.

The most compelling figures in the EUSIPCO paper on Fractional Superlets (Figure 2 being one of them) show that the weighted geometric average implementation of FSLT (Equation 3) show much more cursive transitions between the integer orders. Figure 2 shows a comparison between STFT, CWT, ASLT and FASLT on EEG data with rich time-frequency content. One can easily see that the FASLT provides the best possible representation, for the reason that it is not diluted (like the STFT), not redundant in the upper frequency range (like the CWT) and does not suffer from banding (like the ASLT).

### 3.2.3 Superlets vs. competition - real data

The most important comparison to make is between the SLT and other established methods in how they perform on real neuroscience-specific data. This poses a problem, however, because the ground truth is not known. This calls for an ML-assisted algorithm that can estimate the information content in these power spectra. A summary of this algorithm is discussed in the methods section.

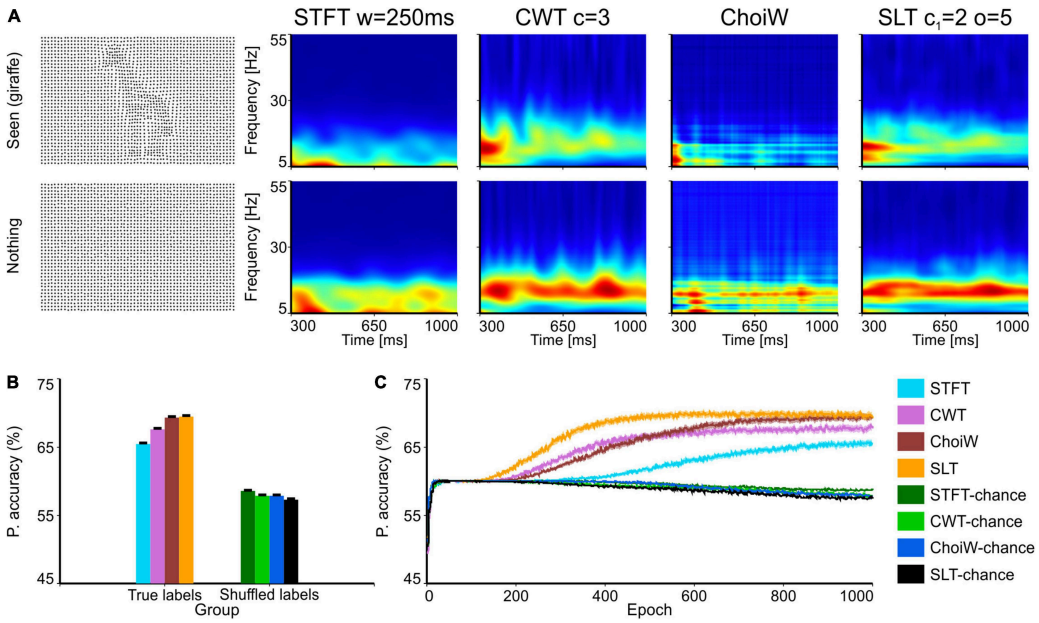


**Figure 4: Performance evaluation results for the EEG data. A.** Two conditions were used: seen (top) and nothing (bottom). **B.** The performance (accuracy) measures for the test (left group) and control (shuffled labels, right group) data sets. **C.** Learning curves for the datasets in B. The data in B and C represents averages over 100 splits. The errors are represented in SEM. Taken from Bârzan *et al.* 2022.

The methods section also introduces the newly developed joint feature permutation technique. The results for the same EEG data as in Figure 4 are presented in Figure 5.

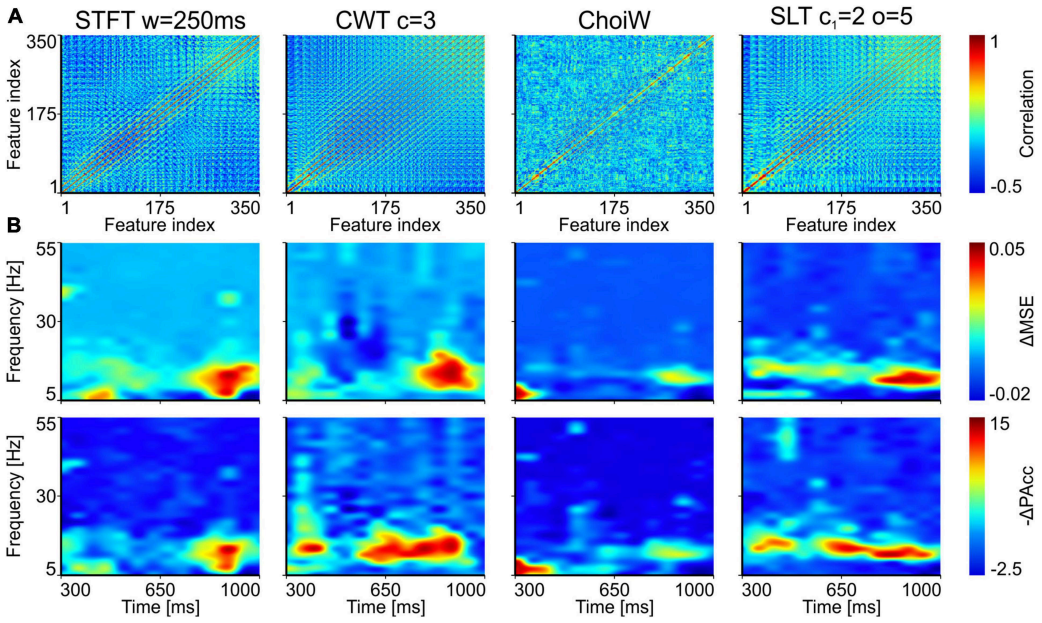


**Figure 5: Joint feature permutation results. A.** The correlation matrices computed for each data set. **B.** JFP results for the difference in MSE ($\delta MSE$) and PAcc ($-\delta PAcc$, in %). Larger values in $\delta MSE$ and $-\delta PAcc$ represent more important features. Taken from Bârzan *et al.* 2022.

### 3.2.4 Discussion on superlets

It is evident from the results in Figure 2 that the fractional SLT is a very powerful tool for working with neural time series data (particularly EEG and LFPs), especially when high frequency oscillations (i.e. gamma oscillations) are the focus. The adaptiveness of the SLT makes viewing large frequency ranges an easy endeavor, due to the fact that it does not suffer from dilution or redundancy. The SLT also has an edge when its information content is measured empirically *via* machine learning, with its representations managing to capture more information from neural data than other time-frequency methods (Figures 2 and 4 are good examples on EEG data).

The popularity of superlets has also taken off since publishing the first superlet paper in 2021. Currently there are 41 papers citing the superlet papers, and people are frequently making inquiries to our lab for additional clarifications and support with using superlets for their data. The range of applications seems to be very large - neuroscience (its intended audience), power electronics, internal combustion engine design, ecoacoustics, radar, material science, cardiology, civil engineering and even astrophysics papers are citing the papers at the time of writing. The interaction with some of these laboratories also spurred more research into the various aspects of superlets, such as creation of a complex input-real output FASLT for radar applications.

## 3.3 GENUS therapy

As previously discussed, superlets shine when looking at high-frequency oscillations such as gamma oscillations (30-80 Hz). Therefore it makes practical sense to try out the superlets in one of the most interesting therapeutic approaches involving gamma oscillations as treatment for Alzheimer's disease. *Gamma ENtrainment Using sensory Stimuli* (GENUS) is a novel therapeutic approach developed by a team of researchers led by Dr. Li-Huei Tsai, which involves administering flickering stimulation in the form of light and sound, particularly at 40 Hz. This

**Figure 6: Still of a gaming session using the BCI** featuring the thesis' author. The depicted game, Audiosurf™, is an action game featuring a spaceship on a track, where the user must avoid obstacles and collect blocks (points). The cadence at which both arrive is given by the agressiveness of the rhythm of the music.

method has been shown to cause reduced genetic expression of amyloid-$\beta$, while also causing microglial phagocytosis of already present amyloid-$\beta$plaques.

We created an experiment in which we measured the gamma entrainment elicited by stimulation with flickering light, in keeping with the GENUS protocol. Our question revolved around which color of light produces the most gamma entrainment. By trial and error we eliminated the red and green colors and were left with only the white (used by Li-Huei Tsai and her team in the original GENUS papers) and blue light.

Results for 3 mice show that, at 40 Hz (the frequency specified in the GENUS papers), blue light produces significantly more gamma power than white light. While further analysis is needed to quantify other parameters (such as the length of persistent entrainment), these results have the potential to change how GENUS therapy is administered.

## 3.4 Real-time brain computer interface

Lastly, we deal with the results obtained using our novel gamma-based brain computer interface described in the methods section. Inspired by alpha and mu rhythm-based motor imagery BCIs, this new implementation manages to do something that is quite rare in the motor imagery BCI field - real time actuation. The Audiosurf™gaming session featured in Figure 6 is a good example of the real time capabilites of the BCI. The game requires good coordination to navigate the track and avoid the incoming obstacles.

Of course, we also needed to quantify the performance of the BCI from an information theoretic viewpoint, in order to rank it against similar implementations. Therefore, we used Shannon entropy to evaluate the information *bandwidth* of the BCI, i.e. the number of bits of information it can convey per unit of time (minutes, seconds). We devised a special task for this, involving producing a gamma power differential across the left and right motor cortices, as one would to control the game, but instead this power differential is used to select from a range of fields (see Figure 7). Each field is a symbol (in the framework of Shannon entropy), and the task is to select the indicated field as quickly as possible. Information transfer rate (ITR) is measured using the following formula:


**Figure 7: User interface for the bandwidth measurement task.**

$$ITR_{BCI} = r \cdot \left( \log_2 N + P \log_2 P + (1-P) \log_2 \frac{1-P}{N-1} \right) \tag{5}$$

where,

- $N$ is the number of fields (symbols),
- $P$ is the probability that the outcome (selected field) is the desired one (highlighted field) when the command is computed,
- $r$ is the polling rate of the BCI, measured in commands per second or Hz.

The ITR results for this field selection task clock the BCI at around 88 bits/minute on average (1.46 bits/second), with an average task performance of 90% (meaning the correct field was selected on 90% of the trials within the alloted time limit). However the highest measured ITR was around 149.4 bits/minute (2.49 bits/second). The theoretical maximum for this particular experiment with 5 fields (as shown in Figure 7) is 200 bits per minute (in each trial the first command correctly selects the indicated field).

These results place our BCI at a particular advantage, given that most motor imagery-based BCIs around 12 bits per minute. There are also visual-based BCIs using steady-state visually evoked potentials (SSVEPs), which rely on stimuli with different flickering frequencies on a computer screen, with the user picking a stimulus by focusing their gaze on them. The flicker rate is then detected by the BCI, making out which object was selected. The maximum bits/minute for a SSVEP-based BCI speller was 325 bits/minute (5.4 bits/second). Note, however, how similar this BCI is to using an eye tracker, a device which measures the position of the pupils and can compute where the user is looking on a computer monitor. In comparison, the gamma-based BCI presented in this thesis does not require the user have sight, given that decent feedback is provided to the user through other sensory modalities.

# 4 Conclusion

This thesis introduces a number of techniques relating to the investigation of fast cortical rhythms (with a particular focus on gamma oscillations) and using them in therapeutical and practical applications. Of particular note is the creation of the Superlet Transform, a robust time-frequency analysis method with a focus on high resolution representations. This new technique has aided us in the creation of the brain-computer interface and also provided good insight into the parameters of gamma entrainment in GENUS therapy. The SLT transcends the boundaries of this thesis, however, with a fair number of researchers employing the technique for their own ends in a wide swath of scientific fields.

The investigation into GENUS therapy yielded a potentially significant result with regards to how this therapy is applied, with blue light flickering apparently being significantly more capable in producing gamma oscillations than the white light used in the original protocol.

Lastly, the real-time BCI introduced in this thesis is also of high scientific and engineering value, due to the novelty factor of using gamma oscillations as control signals non-invasively, using EEG. The novelty lies not only in this peculiar control signal, but also in the responsiveness we achieved with it, tested at a command rate of over 4 commands/second. The information transfer rate (i.e. the bandwidth) of this BCI also places it in the top echelons, with its average bandwidth strongly exceeding other BCIs of this type.

This thesis serves as proof of the highly interdisciplinary nature of neuroscience and is an example of how, with the right tools, fundamental neuroscience concepts such as gamma oscillations can be exploited in various therapeutic and practical applications.