How should I re-reference my intracranial EEG data?

George M. Parish[1], Sebastian Michelmann[2,3], Simon Hanslmayr[1,4]

1: School of Psychology, Centre for Human Brain Health, University of Birmingham, Birmingham, United Kingdom

2: Department of Psychology, Princeton University, Princeton, NJ, USA

3: Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA

4: School of Neuroscience and Psychology, Centre for Cognitive Neuroimaging, University of Glasgow, Glasgow, United Kingdom

Abstract

The analysis of intracranial electrophysiological recordings requires processing choices. Electrical signals are recorded relative to a reference and the choice of that online reference may be sub-optimal depending on the goal of the subsequent analysis. Therefore, a secondary re-referencing operation is often undertaken aiming to increase the signal-to-noise ratio, which can entail transforming the signal in relation to a specific hypothesis. However, comparative studies on this much understudied issue of re-referencing are sparse, which can lead to habitual and ill-informed decision making. This Chapter starts with giving a non-exhaustive overview over common re-referencing schemes before presenting three studies that explore what re-referencing means for cortical alpha & gamma power during a motor task as well as lower frequency power in the medial temporal lobe during a memory task. By revealing how different strategies lead to different observations in the iEEG signal and their modulation by task or behaviour, we demonstrate how significant this early transformative decision is for further analyses.

1.1. Why do we need a reference?

Intracranial EEG offers the exciting prospect of measuring brain activity at high spatial *and* high temporal resolution, which is not possible with non-invasive methods where a researcher needs to compromise on either of the two dimensions. The motivation of an iEEG study therefore often is to utilize this unique strength in temporal and spatial resolution in an experiment. However, there is a fundamental biophysical reality that any researcher wishing to use iEEG faces, which naturally comes with the type of signal we record. This signal represents fluctuations in voltages over time which are caused by weak electrical fields generated by the summated activity of neurons [2]. An iEEG recording at a given channel therefore refers to the dynamic measurement of potential differences, or differences in "electrical pressure", between two points. This can be written as a simple subtraction formula:

$$V_{(t)} = V1_{(t)} - V2_{(t)}$$

Where V(t) is the signal at time (t), and V1(t) is the electrical potential at one recording site at time (t) and V2(t) is the electrical potential at another recording site at time (t). This raises a fundamental question, which is "Is a change in V caused by a change in V1, or V2, or by a change in both?". Referencing describes the process by which we initially chose the recording points V1 and V2 to record a signal of interest V, which is in most cases dictated by clinical needs. Re-referencing refers to carrying out this subtraction offline, i.e. after the initial recordings have been made, to isolate a signal of interest based on a particular research question. Re-referencing is therefore one of the first steps in the preprocessing "pipeline" of an iEEG dataset. Because there is no "one-size-fits-all" approach (as will become clear in the remainder of this chapter), this initial step needs to be the result of an informed decision based upon careful considerations. An in-optimal reference may even lead to erroneous conclusions.

At this point you may be asking yourself, why don't we simply choose an electrically inactive site as a reference? Indeed, the very term "reference" implies a source that is non-active against which the activity of an "active" source can be measured. Unfortunately, this is a myth often encountered in EEG research. In such a hypothetical scenario any change in the signal can be interpreted unequivocally to the change in the "active" electrode. The biophysical reality of the brain, however is that there are no electrically inactive sources (see [3] for an excellent in-depth discussion on this issue[1]). While there are areas in the brain that are electrically more or less active, like grey matter (electrically more active) and white matter (electrically less active), there will be no area in the brain, or on the scalp that is electrically inactive [4]. Even electrodes placed on the bone or on the skin outside the brain will pick up electrical activity that is volume conducted to that site and therefore introduce this activity into recordings if being used as reference.

Let us consider a specific example of a recording where electrodes in the hippocampus are referenced against a scalp electrode placed on the mastoid (i.e., a point at the back of the head behind the ear, which is a popular choice for referencing [5]). After analysing the data, the researcher may find alpha power decreases that are modulated by the task and concludes that alpha power modulations in the hippocampus reflect a specific cognitive process. However, it may well be that the hippocampus itself does actually not show any modulations in alpha power at all, and that instead these alpha modulations are solely introduced by the mastoid reference which happens to pick up alpha signals volume conducted to the scalp. Such an example is demonstrated in [6] (Figure 2) and illustrates one possibility where a wrongful conclusion is made due to an in-optimal referencing choice[2].

So far, we have only considered electrical fluctuations that are generated by the brain. However, several other sources also give rise to changes in electrical potentials that may be picked up by intracranial electrodes, or introduced by reference electrodes. Muscle activity induced by head movements, speech, and chewing for instance; another example is power line noise introduced by electrical sources near the patient (i.e., power sockets, patient bed, medical devices near the patient, etc.). These signals can be amplified by sub-optimal reference choices and, in extreme cases, may render a whole recording unfit for analysis. On

---

[1] This book chapter is mostly concerned with scalp EEG but many of the fundamental biophysical properties also apply to intracranial EEG.

[2] This is not to say that the hippocampus does not show alpha oscillations that are modulated by a task. Indeed several studies suggest that the hippocampus expresses genuine alpha oscillations that are modulated by memory processes (e.g. 7. Staresina, B.P., et al., *Hippocampal pattern completion is linked to gamma power increases and alpha power decreases during recollection.* Elife, 2016. **5**.).

the other hand, re-referencing can be a powerful tool to separate out such artefacts and may even salvage a “lost” dataset.

To summarise, we have clarified that any intracranial EEG researcher is forced to make a decision on how to re-reference the data. This decision has the potential to improve the data quality or to make it worse, and in extreme examples lead to wrongful conclusions. Because no electrically inactive source exists there is currently no gold standard in the field as to how to re-reference an iEEG dataset. Hence the answer to the title question of this chapter *“How should I re-reference my iEEG data”* will be different on a case-by-case basis, depending on the research questions and what the signal of interest is.

The aim of this book chapter is to introduce the reader into the complex issue of re-referencing by providing a (non-exhaustive) overview over the different most used referencing schemes. A secondary aim is to give the reader an intuition about the advantages/disadvantages of different referencing schemes by reviewing the results of two recent empirical studies exploring the effects of re-referencing on a signal of interest. Inevitably, this chapter will raise more questions than it will answer because our current knowledge on this issue is far from being complete. Our hope is to stimulate further research, and to give the reader a few tools to make a better-informed decision on which re-referencing scheme to choose for their dataset.

1.2. How can we re-reference iEEG data

The following (non-exhaustive) list of commonly used re-referencing approaches may guide the reader through the literature on this topic and help them to make an informed decision about the optimal choice of reference for the analyses they pursue (see Figure 1). Before describing these different referencing schemes it is important to define the terminology to avoid confusion. The term ***electrode*** is used to refer to the device that is inserted into the brain, or placed onto the brain which can be either a depth, strip or grid electrode[3]. The term ***contact*** refers to a singular location on that device where electrical contact with the brain tissue is made. The term ***channel*** will be used to refer to the recorded signal which is impacted by given referencing or re-referencing procedures.

[3] In some occasions we also use electrode to refer to recording sites used in scalp EEG.

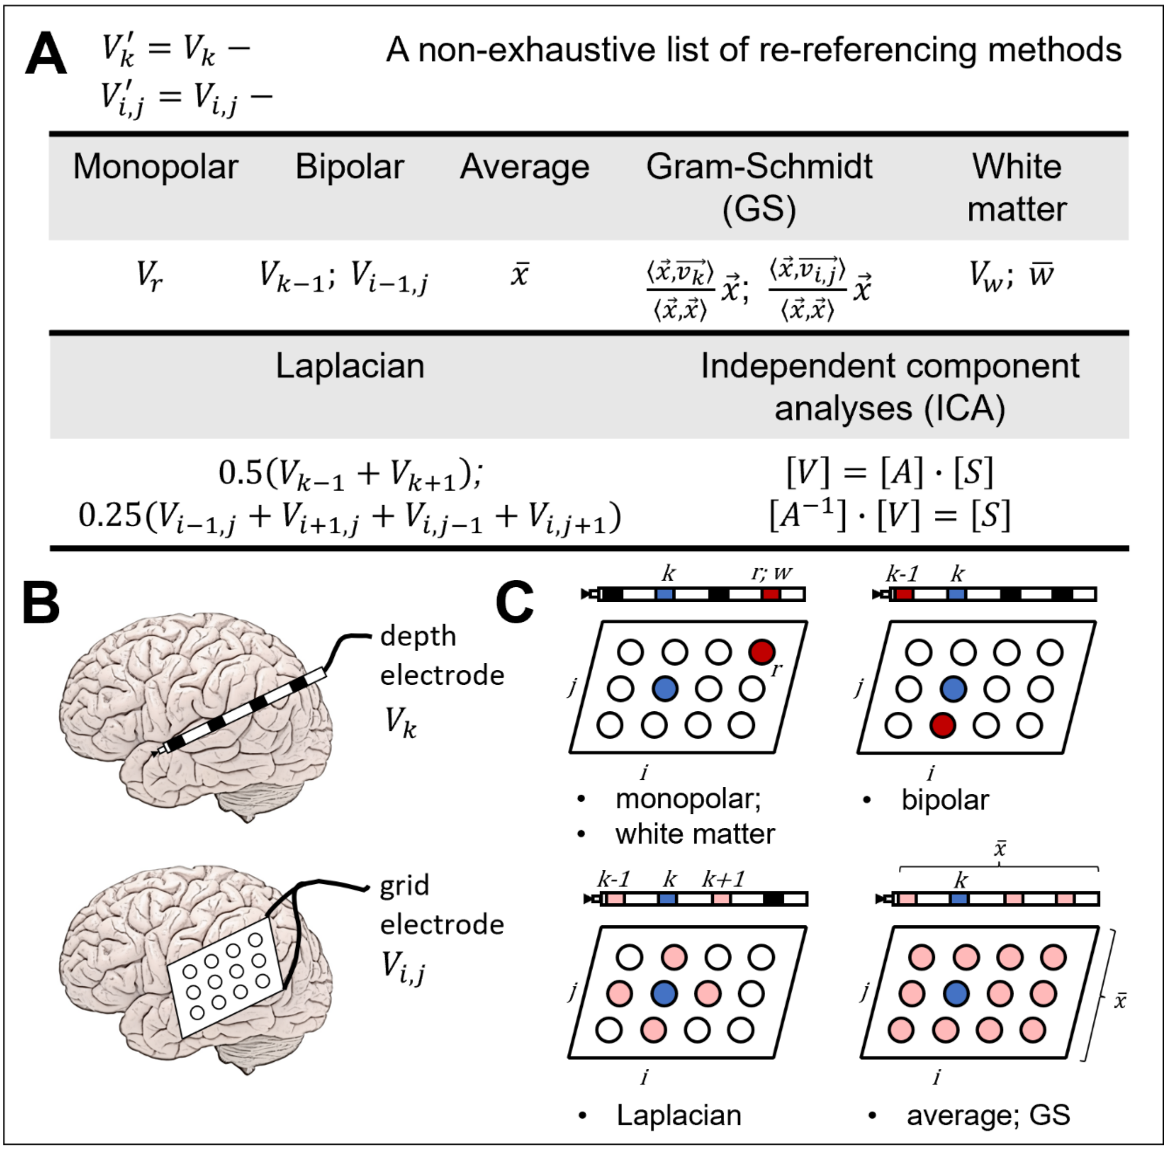

Figure 1: A non-exhaustive list of commonly used re-referencing methods. A: Equations for several re-referencing methods that can be applied to intracranial EEG data, for both 1-dimensional depth electrodes (Vk), 1-dimensional strip electrodes and 2-dimensional grid electrodes (Vij). B: Visualisation of a depth electrode (top panel) and grid electrodes (bottom panel). C: Visualisations of the application of the equations in A to the actual data. Blue nodes indicate the contact in question (Vk or Vij), red nodes indicate single target re-reference contacts and light red indicate a set of re-reference contacts that will be averaged over. Note that these visualisations do not depict possible re-reference choices that are exterior to the brain, for example in the case of the monopolar method, nor do they depict the possibility to average over all depth or grid electrodes within the brain, for example in the case of the average or Gram-Schmidt methods.

1.2.1. Monopolar reference

In intracranial and classical EEG studies, monopolar reference describes the referencing of all contacts against a single contact [2]. This is the typical scenario that is encountered in raw iEEG data, where data may be referenced to a subdural contact, or to a contact that is placed in the bone or on the mastoid [7, 8]. Notably, a researcher's preferred choice of such a reference may not correspond to the optimal choice in a clinical setting.
In the re-referencing step, the researcher can easily change the monopolar reference, by subtracting the newly selected reference channel from all other channels, i.e., at a given channel k that currently reflects the voltage difference between *k* and the online reference *o*, activity is recomputed such that

$$V_k' = (V_k - V_r)$$

where r is the index of the new reference channel. Because both V_k and V_r are capturing the voltage difference to the online reference *o*, the activity from o will cancel out during subtraction. Note that the removal of shared activity that stems from the online reference is a common goal that is shared between many re-referencing schemes.

1.2.2. Bipolar reference

Bipolar re-referencing is one of the most applied re-referencing schemes in intracranial EEG. It describes the subtraction of each channel from its neighbour. Specifically, a new channel k is computed as

$$V_k' = (V_k - V_{k-1})$$

A key goal of bipolar referencing is to highlight activity that is local: by subtracting a channel from its neighbour (typically on the same electrode), activity that is shared between the two channels will cancel out. Bipolar re-referencing can also be described as the spatial derivative; after re-referencing a channel captures the change in activity from one contact to the next [4, 9]. Because activity recorded from the reference contact should affect neighbouring channels to a similar extent, noise and activity from the online reference will cancel out with the bipolar referencing operation. Likewise, however, bipolar re-referencing removes any signal that is shared between two neighbouring channels, which disproportionally affects the low-frequency range. Depending on how much signal and noise is shared on two neighbouring electrodes, bipolar re-referencing can therefore either increase or decrease the signal to noise ratio on a channel (see: [6]). Another key disadvantage of bipolar re-referencing is the inevitable loss of information due to the reduced dimensionality of the re-referenced data: after re-referencing on, for instance, a depth electrode with 8 contacts, activity will be pulled together into 7 channels that reflect the difference between neighbours. The measured activity on the original 8 channels, however, cannot be reconstructed anymore by the linear combination of the new 7 channels, unless there were already linear dependencies beforehand (because the new data has a reduced rank, [6]).

1.2.3. Laplacian reference

Laplacian reference pursues a similar goal to bipolar re-referencing, however, instead of considering a single neighbour, each channel is re-referenced to the average of both neighbouring channels for shaft and strip electrodes, or to the average of its 4 nearest neighbours for grid electrodes [4]. Specifically, for depth and strip electrodes a new channel k is computed as

$$V_k' = V_k - 0.5(V_{k-1} + V_{k+1})$$

On grid electrodes, the new channel at the 2-D position (i,j), can be computed as

$$V_{i,j}' = V_{i,j} - 0.25(V_{i-1,j} + V_{i+1,j} + V_{i,j-1} + V_{i,j+1})$$

The advantages and disadvantages of Laplacian re-referencing are comparable to bipolar re-referencing, however, because all neighbours are considered the researcher does not have to decide on a direction for the operation. A key disadvantage is again the inevitable loss of information due to the reduced dimensionality of the data (see above). This problem is even exacerbated with Laplacian re-referencing, the resulting data will be reduced in rank by 2 for each electrode [6].

1.2.4. Common Average Reference (CAR) and Median Reference

Average reference is a popular reference for scalp EEG. In scalp EEG a coverage of the head can be approximated as a sphere if sufficient (and approximately equidistant) electrode coverage is given. The sum of electrical potentials that are measured on opposite sides of the head should therefore – at least in theory – be zero. Therefore, in scalp EEG, a reasonable approximation of removing activity from the online reference (often mastoid, or Cz) can be achieved by subtracting the average [3, 10].

The broad electrode coverage to approximate the head as a sphere is almost certainly never achieved with intracranial EEG, where coverage is entirely determined by clinical considerations. Furthermore, intracranial contacts measure signal that is more local than the summed potential that is picked up by scalp EEG contacts [11] and depending on the location of a contact, large differences in electrical potential can be observed; the reader may consider, for instance, the difference in amplitude between contacts located in the Hippocampus and in nearby white matter.

Nonetheless, the average across all channels may entail an approximation of activity at the reference, especially if the online reference is very noisy and therefore accounts for a large proportion of the variance.

With average reference, a new channel k is computed as

$$V_k' = (V_k - \bar{x})$$

where $\overline{x}$ represents the average across all channels. Average re-referencing has the advantage of preserving information, it only reduces the total rank of the data by one (because the sum of all channels is 0, each channel is a linear combination of all others; it can be re-written as their negative sum). This advantage, however, may (but doesn't have to) come at the cost of a lower signal to noise ratio in the re-referenced channels and of a potential mislocalization of effects. Specifically, the average across all channels is sensitive to extreme values, e.g., sharp high-amplitude noise on single channels; subtracting the average from otherwise clean channels, may therefore reduce the signal to noise ratio on that channel. Furthermore, the average may be sensitive to high amplitude oscillations from subcortical structures. CAR can therefore lead the researcher to attribute effects to structures that are not involved in the measured neural process. On channels that strongly contribute to the average (e.g., high amplitude channels in the hippocampus), the signal will also appear attenuated (by a factor of 1/N, where N refers to the number of channels) after re-referencing.

The use of a common median reference is an attempt to alleviate the sensitivity of CAR to extreme values. Again, a new channel k is computed as ($V'_k = V_k - \overline{x}$), where $\overline{x}$ now represents the median across all channels, e.g., [12].

1.2.5. Gram-Schmidt orthogonalization

A potential issue with re-referencing is that the subtraction of the new reference can introduce artifacts that were not present on a channel before the re-referencing step (see above). A recent approach to address this issue is the use of orthonormalization between a channel and the average across all other channels via the Gram-Schmidt process [13]. Specifically, the new channel k is computed as:

$$V_k' = V_k - \frac{\langle \vec{x}, \overrightarrow{v_k} \rangle}{\langle \vec{x}, \vec{x} \rangle} \vec{x}$$

where $\overrightarrow{v_k}$ represents the time-series of channel k, and $\vec{x}$ represents the time series of the average across all channels except k; the brackets denote the dot product. This procedure effectively removes the part of the signal at V_k that is shared with the average across all other channels, reducing the risk of introducing artefacts from other channels. As the average is scaled by the inner product of the channel and channel average, this operation can boost smaller amplitude local signal whilst bringing down higher amplitude global signal.

1.2.6. White matter reference

The idea behind the use of a white matter reference is to select a single channel or the average of a group of channels that pick up little to no signal. Because electric potentials that are associated with neural activity are generated in gray matter, contacts located in white matter are assumed to not pick up signal and only reflect shared noise, e.g., from the online reference.

The new re-referenced channel is then computed as

$$V_k' = (V_k - V_w)$$

where w is the index of the selected white matter channel or

$$V_k' = (V_k - \overline{w})$$

where $\overline{w}$ is the average activity across a group of the white matter channels [14, 15]. Importantly, the assumption of white matter contacts being "silent" is overly simplified; a recent study by [4] demonstrated that contacts in the white matter record a mixture of (zero lag) volume conducted signal from nearby gray matter and surprisingly from distant gray matter, i.e., signal that was carried by the white matter fibers themselves [4].

Box: Identification of white matter contacts in intracranial EEG

One of the challenges for the use of white matter reference is to determine which contacts are suitable. Contacts are often located via MRI scans from before and after surgical implantation of electrodes, that are then normalized against one another – a transformation that does not always give perfect anatomical information. To quantify the amount of white/gray matter surrounding a contact, Mercier et al (2017) propose a Proximal Tissue Density index (PTD) based on white/gray matter estimates from a Freesurfer parcellation. PTD is computed within a 3x3x3mm cube around the centroid of an electrode contact as PTD = (VoxGray - VoxWhite) / (VoxGray + VoxWhite), where VoxGray and VoxWhite describe the number of gray and white matter voxels respectively. This measure results in a value between (-1) for contacts that are surrounded entirely by white matter and (+1) for contacts that are surrounded entirely by gray matter. When selecting white matter contacts for re-referencing, it is also useful to assess properties of the recorded signal.

Parish et al (2022) propose an approach that is informed by electrophysiological signal properties. They reason that, as a 'silent' channel, a white matter contact should have a lower signal to noise ratio than other channels. They demonstrate that a white matter contact is characterized by less variability in the local field potential (LFP), and by less low frequency (2-12Hz) power and gamma frequency (60-90Hz) power than contacts that are situated in gray matter (e.g., in MTL gray matter). The authors propose to perform white matter channel-selection based on a threshold, where, for instance, WM theta power should not exceed 20% of that for MTL theta power (compare Figure 2; note that contacts in the hippocampus typically express stronger signal than other channels [1]).

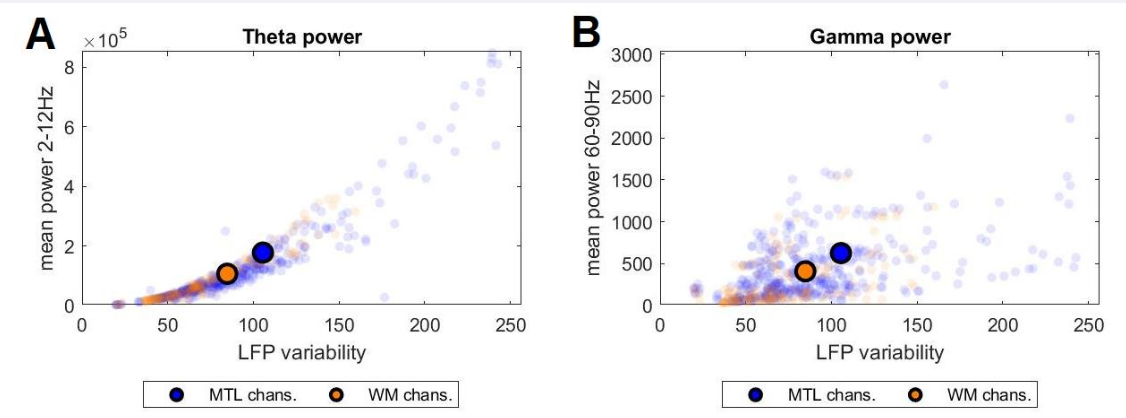

Figure 2: Analyses from an episodic memory paradigm, using iEEG data recorded with depth electrodes, originally referenced using either a bipolar method or a reference contact close to the recording source. **A-B**: Mean power in the theta (**A**; 2-12Hz) and gamma (**B**; 60-90Hz) frequency bands for contacts situated in the medial temporal lobe (MTL; blue) and white matter (WM; orange), against variability of the local field potential (LFP; calculated as 2 x standard deviation). Blue and orange dots indicate individual MTL or WM channels, respectively, whilst blue and orange circles indicate the centroid of all MTL or WM channels, respectively.

1.2.7. Independent Component Analysis

Independent Component Analysis is a method that “unmixes” data into its underlying latent components [16]. It is widely used in EEG to subtract out activity that stems from ocular movements and blinks, but also electrical noise [17]. The reasoning to apply ICA as a re-referencing method, is that neural activity that is measured by a given channel reflects a mixture of local activity, online reference activity, electrical noise and volume conducted activity from nearby areas, but the researcher is blind to the nature of that superposition of activity. ICA can unmix the recorded data into its underlying independent sources (note that noise should be independent from neurally generated activity). The decomposition of the data into its underlying independent components can then be used to systematically eliminate certain components from the data. Because a linear combination of independent components leads back to the original data, it is possible to inspect how much each component affects each channel. Crucially, activity and noise from the online reference mixes into all channels to a similar extent because recorded channels reflect the voltage difference between their respective contact and the online reference contact. It follows that a component that captures activity from the online reference is very global. Discarding this global component should therefore eliminate undesired parts of the data, while leaving the rest of the data intact [6, 18].

This reference can be thought of as a data-driven re-referencing scheme because the coefficients of the computation are learned from the statistical properties of the data at hand. While average reference, for instance, assumes that the average is a good approximation of broad noise and noise at the online reference (using the same coefficient 1/N for every channel), ICA learns coefficients that optimally isolate components that are independent from the rest of the data. A spatially broad component can then be identified by the researcher and be removed from the data.

1.2.8. Spatio-spectral decomposition and tailored spatial filtering approaches

The example of ICA illustrates that it is useful to think of the re-referencing step as spatial filtering operation which highlights certain properties of the data. The bipolar re-referencing of a channel, for instance, is a spatial filter with the fixed coefficients [1, -1] on neighbouring contacts (and zero otherwise), whereas ICA learns coefficients to extract each component from the data. Consequently, the researcher might wonder whether tailored spatial filters should be applied directly for the purpose of highlighting desired properties of the signal. One recent such application is the use of spatio-spectral decomposition (SSD) [19] for intracranial EEG data [20].

SSD is a spatial filter that operates on the spectrum of the time series, to maximize power in a selected frequency band over its flanking frequencies. If the goal of the analysis is to extract specific oscillations with high signal to noise ratio, the researcher may therefore opt to use this method directly in lieu of re-referencing. Similarly, it is noteworthy that other spatial filtering methods (e.g., logistic regression, or linear discriminant are often used for classification; [12, 21]) can learn coefficients including those afforded by common re-referencing schemes. It may therefore be preferable to learn spatial filters on the data directly for analysis purposes, without an intermediate re-referencing step that could potentially result in a loss of information (see above).

2. What does re-referencing do to my data? An empirical comparison between different re-referencing schemes

In this section, we collect results from the relatively sparse number of studies that compare how your re-referencing choice reflects on your data. These studies work with data collected from two very different experiments: one an associative memory paradigm that is interested in low frequency theta oscillations in the hippocampus, the other a gesture decoding paradigm that is interested in mid to high frequency alpha and broadband gamma oscillations in the motor cortex. The relative differences in outcome indicate that your choice of re-referencing method should be part of a wider strategy that is driven by your hypothesis, as each method will transform your data relative to its own assumptions.

The data compared here is obtained via intracranial stereotaxic EEG (sEEG) obtained by recordings from depth electrodes which are implanted into the brain to target potential focal points of epileptic seizures for monitoring and assessment of surgical feasibility. As these electrodes penetrate the deeper structures of the limbic system, contacts along the length of the electrode offer recordings from lateral contacts which pick up signals from many neuronal groups – near and far. First, we will look at how different re-referencing methods perform as a means of data cleaning, as well as whether re-referencing affects the proportion of channels that are responsive to the task at hand. From there, it is important to ask what is left in your electrical signals after all this pre-processing, and more importantly – what is it telling you? This is described in section 2.2. which illustrates the effect of re-referencing on basic signal metrics such as power in specific frequency bands, and task-related differences between these signals as observed in a memory task and a motor task.

2.1. What is left in your data after re-referencing?

A recent study [14] compared how several re-referencing methods transformed the raw signal, with a specific interest in obtaining local signals for analyses in the broadband gamma and alpha frequency bands. During data acquisition, the data was originally referenced using the mean of two white matter contacts. This study found that the signal was quite often contaminated when applying a monopolar re-reference from white matter areas, whereby task-related activity that was previously localised to the monopolar reference channel was introduced to all other channels to become more global in nature (see Figure 3A-B). This led to more channels being responsive to the task at hand (Figure 3C-D), however, this did not translate into an increase in overall power in these frequency bands (Figure 3E-F). Another important observation made in this study is that neighbour driven methods such as bipolar or Laplacian seem to be good at eliminating global activity from the combined set of macro contacts, leading to a relatively small amount of correlation between channels (see Figure 3A-B). The ability of neighbour driven methods to remove global components and reduce the contamination of task-related activity enables further analyses to focus on the most responsive channels (Figure 3C-D) that contain the most local amount of useful signal (Figure E-F).

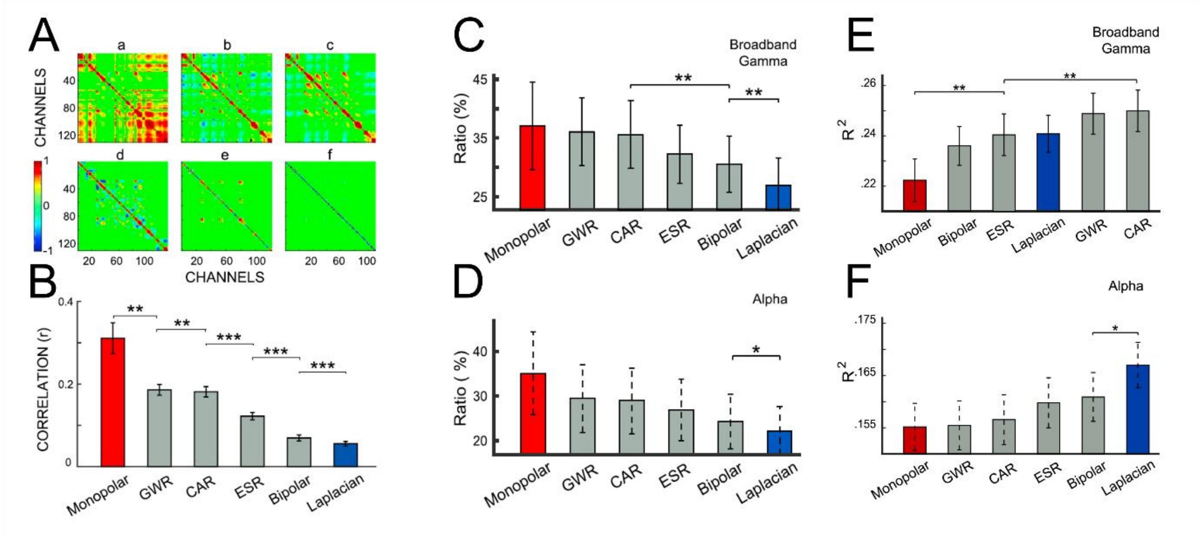

Figure 3: Analyses from a motor task paradigm, using sEEG data originally referenced to white matter contacts. Time-frequency analyses here do not perform a 1/F correction. **A-B:** Signal correlation for different referencing methods. (**A**) Correlation matrix from Subject 12 for the six referencing methods: (a) monopolar; (b) grey/white matter (GWR); (c) common average (CAR); (d) electrode shaft resonance (ESR; where channels are referenced to the average across the entire grid or depth electrode); (e) bipolar; and (f) Laplacian. Colors correspond to the correlation between two specific channels. The correlation between channels varies across the methods. (**B**) Average Pearson's correlation and standard error for the six referencing methods. Asterisks denote the significance of the difference between correlations established using paired t-tests: *** ($p < 0.001$), ** ($p < 0.01$). These statistical results are shown only for the nearest pairs that show a significant difference. **C-D**: Fraction of all channels that are related to the task for different referencing methods. For each subject, we calculated the ratio of task-related channels by dividing the number of task-related channels by the number of all channels. (**C**) Mean (averaged across subjects) and standard error of the ratio of task-related channels for broadband gamma power. (**D**) Mean (averaged across subjects) and standard error of the ratio of task-related channels for alpha power. Asterisks denote the significance of the difference between the ratio of task-related channels for adjacent referencing methods, established using paired t-tests: ** ($p < 0.01$), * ($p < 0.05$). These statistical results are shown only for the nearest pairs that show a significant difference. **E-F**: Coefficient of determination (R2) for different referencing methods, which determines how strongly alpha or broadband gamma power was modulated by the task [22-24]. (**E**) Mean and standard error of R2 for broadband gamma power, calculated across all channels from all subjects. (**F**) Mean and standard error of R2 for alpha power. Asterisks denote significance of the difference (paired t-test) between R2 values for referencing methods: ** ($p < 0.01$), * (p < 0:05). These statistical results are shown only for the nearest pairs that show a significant difference.

Neighbour driven re-referencing methods are therefore commonly applied to maximise local signals. However, [6] demonstrate that ICA may be better suited for this goal (see Figure 4).

The authors compare the performance of ICA in extracting local signal, to the performance of bipolar referencing. Simulations suggest that ICA outperforms bipolar referencing in sensitivity (i.e., in isolating signal from local sources) and specificity (discarding activity from distant sources). Indeed, bipolar reference only performed reasonably well, when signal was very local and noise levels were low, however, ICA still performed better under these conditions.

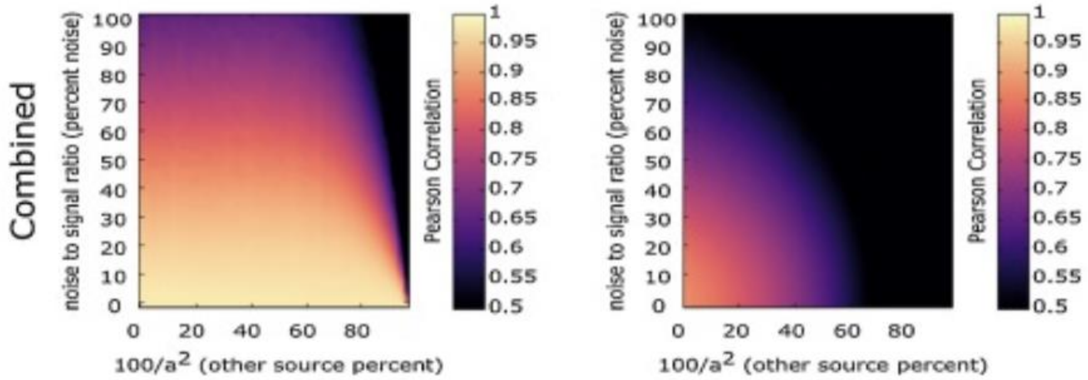

Figure 4: Combined correlation (sensitivity*specificity) of extracted ICA components (left) and bipolar channel (right) with simulated sources. After re-referencing, the local component and the bipolar channel were correlated with an underlying simulated source (sensitivity) and with an interfering distant source (1-correlation, i.e., specificity) under various levels of noise (y-axis) and spatial mixing (x-axis; mixing from a distant source before re-referencing/unmixing, where $1/a^2$ is the scaling factor of the source at distance d that decays with $1/a^d$). Bipolar referencing only performed well when signal was very local and noise levels were low, yet ICA outperformed bipolar re-referencing under all conditions.

When choosing a re-referencing method, one thing to consider is how it will interact with the signal in relation to artefacts. Here, we consider an artefact to be an undesirable sharp transient in the grounded electrical signal, the causes of which can be varied. They might be introduced by electrical interference near to the recording or grounding sites, such as 50Hz line noise or similar electrical bursts. Since sEEG depth electrodes are implanted to identify the origin of epileptic seizures, interictal epileptiform discharges (or IEDs) produced by a pathological area are a common occurrence. Whilst these are a desirable observation in a clinical setting, they can distort offline analyses of the electrical signal and understandably lower the patient's performance on memory or attentional tasks (see also Chapter 5). When pre-processing sEEG data, it is often considered important to identify and remove trials with

IEDs and other artefacts. Artefact identification and removal can be done either manually or algorithmically depending on time and the volume of data.

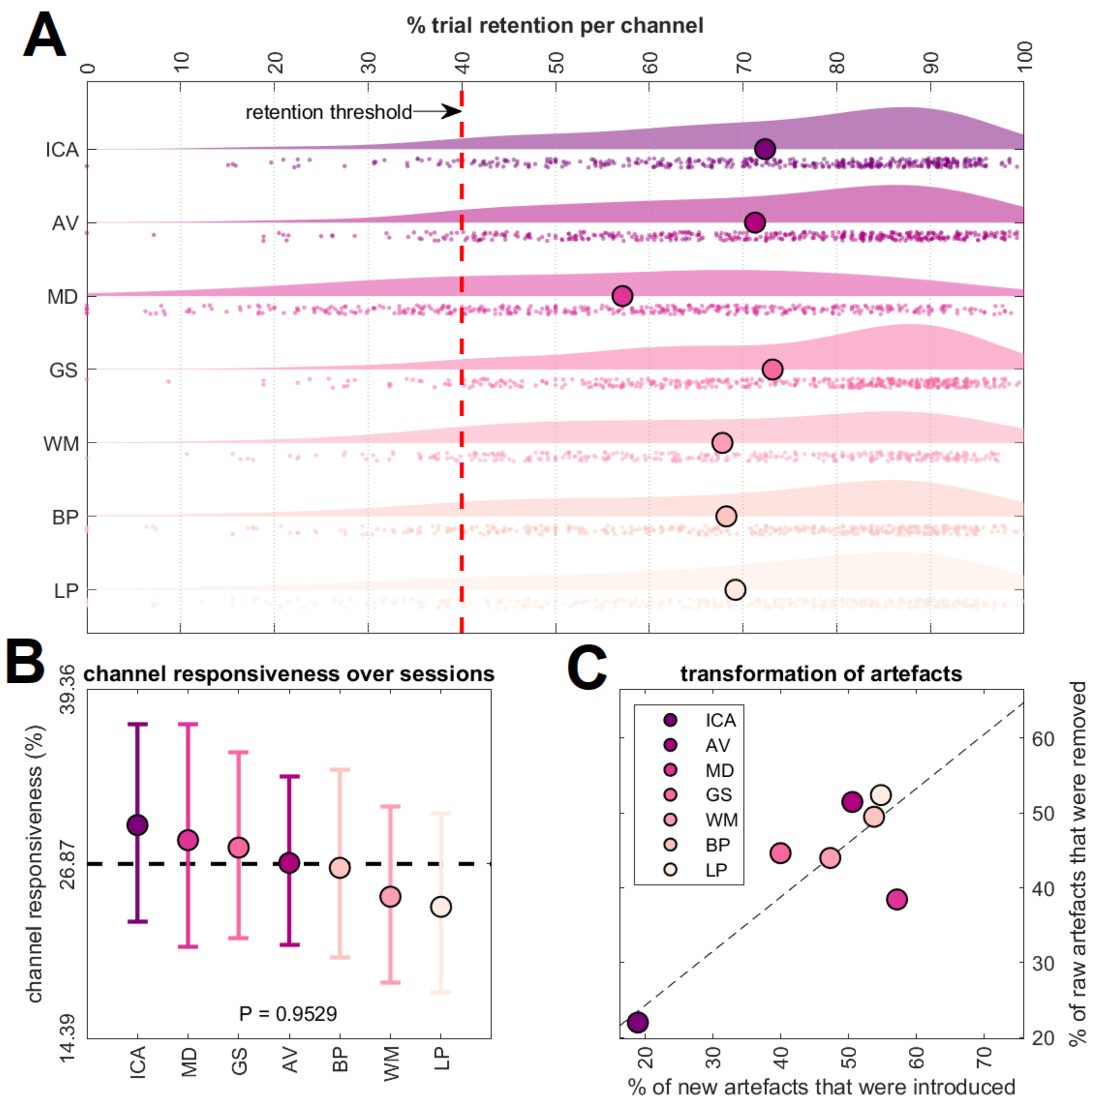

Figure 5: Automatic artefact rejection on a large dataset of 16 subjects, looking at macro contact data from sEEG depth electrodes. If the variance of a trial, calculated as the root mean square (RMS) or Z-score normalisation of the LFP, was above a threshold of 4, then the trial was indexed as artefactual. **A**: visualisation of the retention percentage across channels for 7 commonly used re-referencing methods, independent component analyses (ICA), common average (AV), median (MD), Gram-Schmidt (GS), white matter (WM), bi-polar (BP), and Laplacian (LP). Re-referencing was applied before artefact rejection. The red dotted line indicates the threshold for retaining a channel in the analysis (i.e., 40%). The order of methods is determined by overall theta power (1-12Hz; see Figure 6). **B**: the percentage of channels that showed significant inter-trial phase consistency between 1-20Hz. A 1-way ANOVA indicates that the re-referencing method does not have a significant effect on channel

responsiveness (p = 0.9529). C: proportion of artefacts from the raw data that were no longer present after re-referencing (y-axis), alongside the proportion of artefacts in the re-referenced data that were not present in the raw data (x-axis), essentially: artefacts removed and artefacts introduced by re-referencing method. A line of best fit for both was calculated from information taken from underlying channel datapoints, where scatter points represent averages across re-referencing methods. Figure taken from [14].

[25] employed an associative memory paradigm and compared data retention after applying several commonly used re-referencing methods (described in Figure 1) on sEEG data recorded from the medial temporal lobe (Figure 5). Individual trials were removed after applying an automatic artefact rejection algorithm which was sensitive to sharp transient offsets in the signal. The retention rate of individual channels was then recorded, as well as the proportion of channels that were responsive to the task after an inter-trial phase consistency (ITPC) check between 1-20Hz. To determine the rate by which re-referencing transforms artefacts that are identified in this way (Figure 5C), automatic artefact rejection was applied to identify undesirable trials in both the raw and the re-referenced data. The relative difference in identified artefacts is then described as the number of newly introduced artefacts (i.e., trials that did not have an artefact in the raw signal but now do in the re-referenced signal) and the number of no longer existing artefacts (i.e., trials that did have artefacts in the raw data that are no longer present after re-referencing). In general, artefacts were both removed and introduced at a similar rate, indicating that re-referencing is not an optimal strategy for the removal of sharp transient spikes from the signal and should be used in tandem with more targeted methods. The data presented here is ordered by the amount of power in the theta frequency band (1-12Hz), and as such we can see a slight positive correlation between power and the percentage of responsive channels following an ITPC check – though a 1-way ANOVA suggests that one's choice of re-referencing scheme does not have a significant effect on channel responsiveness.

Whilst this comparative analysis does not equate to any tangible data quality metric, it can still be informative to see how re-referencing methods transform the data. In terms of overall data retention, there is little difference between most re-referencing methods except for the median method (MD), which resulted in the removal of far more trials per channel than any other method. Figure 5C shows that the median method is an outlier in how it transforms artefacts within the data: it fails to sufficiently reduce the amplitude of pre-existing artefacts due to its nature of being less sensitive to outliers. Furthermore, it appears to introduce artefacts to more channels than any other re-referencing scheme, which is surprising given the typical reason one might employ its usage, leaving this method with significantly less data than other methods after automatic artefact rejection. In comparison to the use of the mean average (AV), it appears that the high amplitude outliers of a localised artefact are sufficiently offset within the mean, such that the subtraction of the mean strikes the right balance between reducing the localised artefact without introducing it to the other contacts on the electrode. In contrast, the Gram-Schmidt method seems to preserve potentially useful signals by removing more artefacts than it introduces, resulting in a slightly higher trial retention than most of the other methods. This is likely due to the nature of the Gram-Schmidt method, which subtracts variance that is shared across the signal, making it less likely that more localised artefacts will be introduced to other channels. Most notable however, is the way in which ICA works to clean the data of artefacts. It seemingly transforms the data less than other methods (introducing and removing fewer artefacts), though it equally manages to retain a high proportion of trials in comparison to other methods (Figure 5A) and also the highest number of responsive channels (Figure 5B). This suggests the power of using a data driven method to efficiently identify and remove common noise that exists across a set of independent channels.

In line with [14] (see Figure 3) the Laplacian method here returns a lower proportion of channels that are responsive to the task – perhaps indicating the effect that a neighbour driven re-referencing method can have for producing a more localised signal with a lower proportion of task-related channels and a higher signal to noise ratio. Here, the monopolar method also returns a relatively lower proportion of responsive channels relative to other methods, in contrast to [14], where task-related information contained within the monopolar reference was thought to contaminate other channels. This might indicate the strength of using the data to identify appropriate white matter monopolar contacts for re-referencing (see Figure 2).

2.2. What is your data telling you after re-referencing?

We next consider what the data is telling you after re-referencing, looking at two studies. We first continue with the associative memory task [25], looking at power, peak frequencies and memory effects in lower frequency band, followed by another motor task that compared the effect of re-referencing methods on decoding accuracy in higher frequency bands [26].

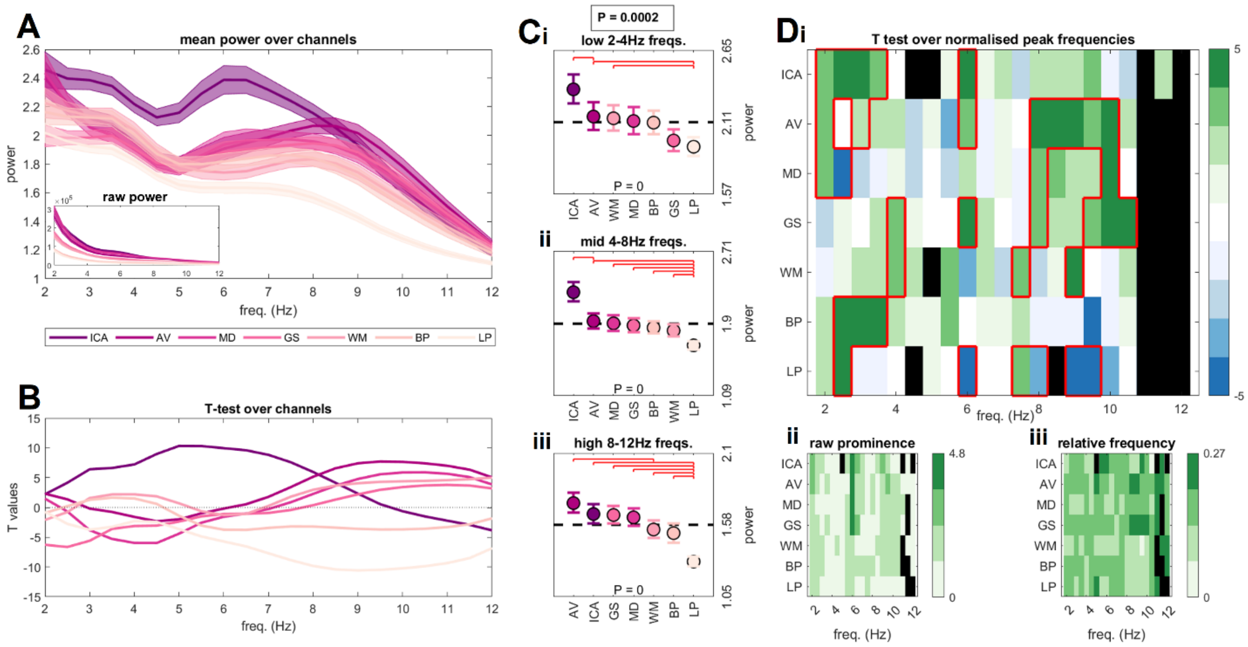

Figure 6: Comparing power and peak frequencies in macro contact data for 7 commonly used re-referencing methods, independent component analyses (ICA), common average (AV), median (MD), Gram-Schmidt (GS), white matter (WM), bi-polar (BP), and Laplacian (LP). The order of methods is rank ordered by overall theta power (1-12Hz). For this analysis we used all available channels regardless of location or responsiveness to the task, except those that were rejected after automatic artefact removal. **A**: Comparative power differences across re-referencing schemes using all available channels per re-referencing method in an independent manner. Power has been 1/F corrected throughout this Figure, though the raw power is shown in an inset of **A. B**: Paired t-test over channels, where each re-referencing scheme was compared to the average of all other re-referencing schemes at every frequency bin. **C**: The mean power across re-referencing schemes was split into three windows between 2-4Hz (**Ci**), 4-8Hz (**Cii**), and 8-12Hz (**Ciii**), rank ordered by magnitude where paired t-tests indicate the nearest rank-ordered neighbour with significantly less power (red lines, p<=0.05). A 2-way ANOVA (top p-value) indicates interaction of re-referencing scheme by frequency and 1-way ANOVAs indicate a main effect per frequency bin (p-values within subplots). **D**: Peak frequency differences. Peak frequencies were detected on a channel-by-channel basis, where the prominence of a peak equates to the difference in power between a local maximum and neighbouring minima (**Dii**). The frequency of detected peak frequencies was also recorded (i.e., the number of times a peak was found at any given frequency bin) (**Diii**), which was normalised across re-referencing methods (such that each frequency bin summed to 1). After applying a 2-sample t-test on the multiplication of prominence and relative frequency, T-values (**Di**) indicate where any given re-referencing

scheme has a higher number of larger peaks than the average of all other re-referencing methods (or vice versa), where significance is indicated by red boxes (p<=0.05, FDR corrected). Black squares in Di indicate insufficient common channels for statistical comparison, black squares in Dii-Diii indicate that no peaks were found.

First, we consider the power of specific frequencies from a signal obtained during an associative memory paradigm [25]. A Wavelet time-frequency decomposition with 1/F correction was applied to the re-referenced signal, where 1/F pink noise was reduced (estimated separately per channel) to enhance the analyses of low frequency oscillations. Low frequency oscillations tend to be more ubiquitous in the brain, as they are important for the long-range synchronisation of neuronal processes, whilst high frequency oscillations tend to be more derived from local neuronal circuits [27]. This means that lower frequencies will return higher amplitudes in the power spectrum, as more neurons are entrained together in phase by more global and low frequency rhythms. In this study the medial temporal lobe is under observation, where low frequencies are thought to be mostly driven by the theta frequency (typically 2-8Hz) which dominates there, rather than the alpha rhythm (typically 8-15Hz) that is dominant in cortical areas [15]. By plotting the average 1/F corrected power across all re-referencing methods (Figure 6A), we can see that macro contacts seem to have two components in the theta frequency range and one component in the alpha range (slow theta at ~2-4Hz and fast theta at 4-8Hz and possibly alpha at 8-12Hz).

Different re-referencing methods have varied effects on each of these low frequency components, which roughly fit in line with their respective assumptions that were described earlier. For example, bi-polar (BP) and Laplacian (LP) both vastly reduce the 4-12Hz components (Figure 6A-B: comparative power across all channels; & 5Cii-iii: statistical comparison of binned mean power for common channels). As low frequency oscillations tend to be a more global phenomenon that is often sampled at many neighbouring recording sites, these signals are therefore reduced here by neighbour driven re-referencing methods, which will result in an emphasis being placed on the more localised faster frequencies as was seen in the previous study, where LP resulted in higher gamma power [14]. Another potential issue with such local BP and LP referencing methods is that they may enhance low frequency activity at tissue borders between areas where oscillations are homogenous within an area but different between areas which would appear as a ring-like shape of enhanced power. This is also reflected in the peak frequency comparison, where LP in particular has significantly fewer and smaller peaks in the 6-10Hz range when compared to other re-referencing methods. In line with previous studies [14], the BP operation does not reduce global components as much as the Laplacian, though interestingly it does return a higher slow theta component than most other methods.

Another method that works in line with its assumptions is the Gram-Schmidt method, which works to remove only the part of the signal that is shared with the average across all channels. Due to the nature of the signal, this typically entails reducing higher amplitude low frequencies whilst enhancing lower amplitude high frequencies (as shown by the upwards trending curve in Figure 6B), which in effect produces a more localised signal. Similarly, other methods that also enhance the putative alpha frequencies (8-12Hz) are the common average (AV), median (MD) and white matter (WM) methods (Figure 6A-C), where AV in particular produces a larger number of prominent peaks in the 8-10Hz range (Figure 6Di). This might suggest that these methods have a tendency to introduce signals from cortical areas where alpha rhythms (8-12Hz) are dominant. By far the method that maintains the most amount of potentially useful theta is ICA. It produces significantly more power in the 2-8Hz frequency bins, indicating that it does not introduce cortical alpha into the hippocampal signal. Similar to the neighbour driven BP and LP methods, ICA also produces a large number of peaks in 2-4 Hz frequency range, though unlike those methods ICA also enhances a 6Hz component which is seemingly lost when applying most other methods.

From this comparative analysis, we might be able to ascertain that the 2-4Hz component is thus the most localised rhythm within the MTL, as it is not significantly reduced by neighbour driven methods such as BP and LP. The 6Hz theta component can therefore be considered a more global theta rhythm as it is almost entirely reduced by these neighbour-driven methods whilst also maximised by the data-driven nature of ICA. Additionally, it might be prudent to caution against the over interpretation of faster theta or alpha frequencies in the MTL if one has applied a more global re-referencing akin to the common average or white matter monopolar, as it might simply have been introduced from cortical areas (this same pattern is also implied by findings from [14]; see Figure 3).

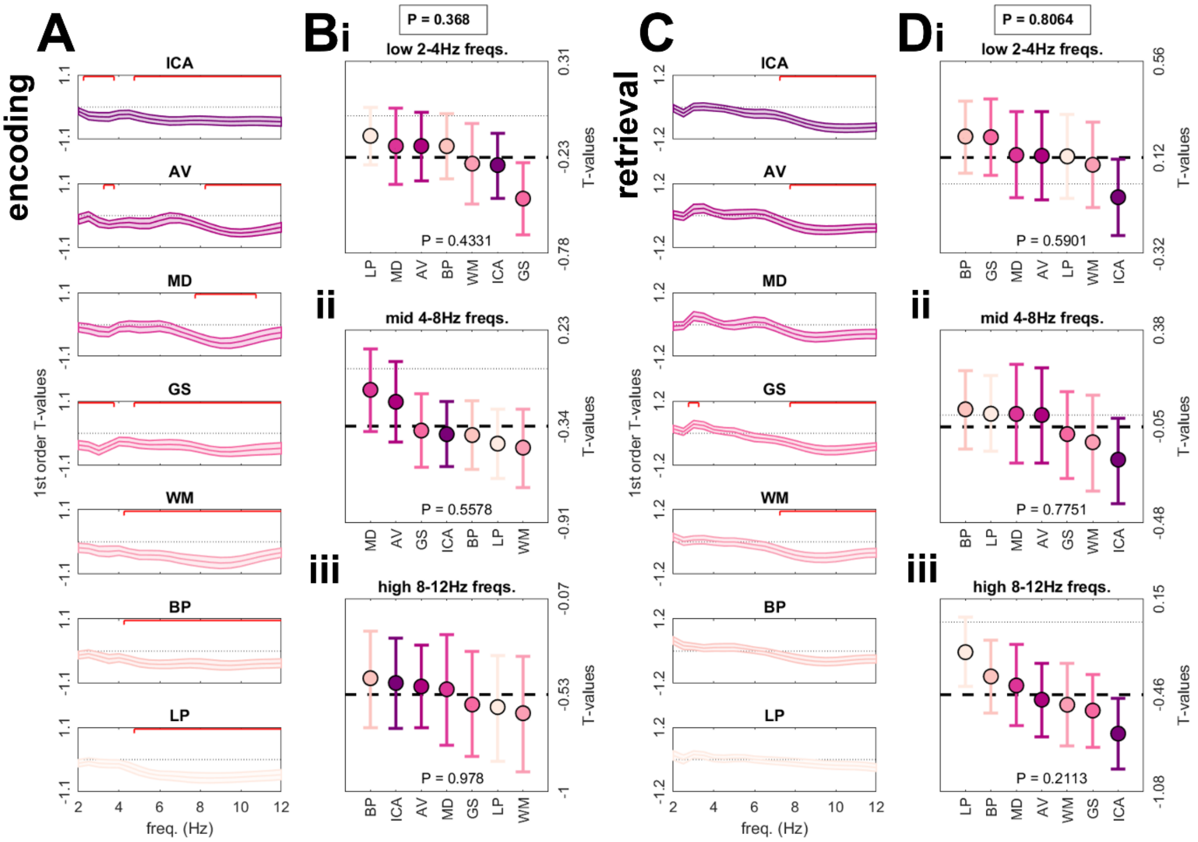

Figure 7: Memory effects for contacts situated in the hippocampus, that indicated significant inter-trial phase consistency (1-20hz). Comparing between 7 commonly used re-referencing methods, independent component analyses (ICA), common average (AV), median (MD), Gram-Schmidt (GS), white matter (WM), bi-polar (BP), and Laplacian (LP). The order of methods is rank ordered by overall theta power (1-12Hz; see Figure 6). Independent t-tests were applied on a channel-by-channel basis for power across hit trials minus miss trials, both for encoding (**A-B**) and retrieval (**C-D**) trials. 1-sample t-test applied to the resultant T-values (**A/C**) to indicate significant differences to zero (red lines, p<=0.05, FDR corrected). T-values were further split into 3 frequency bins 2-4; 4-8 & 8-12Hz (**B/D**), where both a 2-way ANOVA (top p-value) indicates interaction of re-referencing scheme by frequency bin and 1-way ANOVAs indicate the main effect per frequency bin (p-values within subplots).

Next, [25] looked at the effect of re-referencing on memory effects (Figure 7). Data was obtained by way of a cued recall memory paradigm, where subjects were presented with a cue followed by a pair of images (termed the encoding phase). After several trials, and a distractor task, cues were subsequently presented again and the subject indicated how many images they could remember (termed the retrieval phase), before identifying the paired images from a selection on screen. Memory effects refer to the differences between successful memory trials (i.e., hits) where the participant could correctly retrieve the memory, and unsuccessful memory trials (i.e., misses) where the participant could not correctly retrieve the memory. These memory effects were analysed during the memory formation (i.e., encoding) phase of the experiment, and during the retrieval phase of the experiment. Theta oscillations have long been thought to play a formative role for such memory operations [28], where both positive and negative effects (i.e., where hits have more theta power than misses, or vice versa) have been reported. It might be the case that this effect is made up of both a narrowband and broadband component, where contrasts in paradigm, recording technique and referencing scheme might emphasise one of these components over the other [29]. Therefore, it is important to consider whether or not re-referencing has any interaction with such memory effects.

Overall, the data from the study produced large negative fast-theta and alpha (6-12hz) effects – indicating that trials where subjects performed worse had a larger amount of this frequency band, as backed up by the literature [29]. As well as this, there might be a small positive slow-theta (2-4Hz) retrieval memory effect - indicating that trials that were successfully remembered might have higher theta power (see GS data shown in Figure 7C). However, there was no significant interaction between re-referencing method and frequency which suggests that re-referencing did not strongly affect memory effects for the macro contact data. Nevertheless, there were some subtle yet insignificant differences. In general, re-referencing with ICA, WM, AV, MD and GS seemed to produce more negative memory effects in the fast theta and alpha frequency range, particularly for retrieval effects, whereas local referencing schemes such as BP and LP did not show these effects. This may indicate that the memory related alpha power decreases [7] in the MTL originate from a spatially broadly distributed source as opposed to very local sources. Importantly, these negative fast theta and alpha memory effects are unlikely to be introduced by cortical signals as demonstrated by the ICA results. It is interesting to note that the only method to obtain significant theta effects in both memory contrasts was the Gram-Schmidt method, which produced a negative theta effect (hits < misses) at encoding and a positive theta effect at retrieval (hits > misses). It is also worth noting that the Laplacian method completely abolishes any observable effects at retrieval (see Figure 7C), perhaps indicating that a neighbour-driven method reduces too much of the global signal in lower frequencies across both hit and miss data (see Figure 6).

Another study [26] complements these low frequency findings related to memory, by looking at gesture decoding accuracy for a motor task across several frequency bands (see Figure 8). This entailed usage of a classification algorithm that first built a feature vector based upon spectral power (normalised to a baseline period) obtained over overlapping segmentations of the signal, where features constituted the mean power across frequency bins. Next, a set of channels were selected that maximised these features using a search optimisation algorithm, where power was further normalised across trials to eliminate inter-channel differences. A support vector machine was then applied for the classification of multiple hand gestures. This process was repeated in a 10-fold cross-validation process to obtain a comparable gesture decoding accuracy measure for each re-referencing method.

Importantly, it was found that the adoption of any re-referencing method made a significant improvement on decoding accuracy, implying that a sufficient amount of non-task related noise was removed by re-referencing. In general, this study found that the more localised the re-referencing method the higher the decoding accuracy across most frequency bands, likely due to the ability of localised re-referencing methods (such as the neighbour-driven Laplacian and bi-polar) to retain more task-related information by reducing globally shared signals. However, as the spatial topography of implanted electrodes was an essential feature of the decoding algorithm, then it would be expected that re-referencing methods that improve spatial resolution would also improve decoding performance.

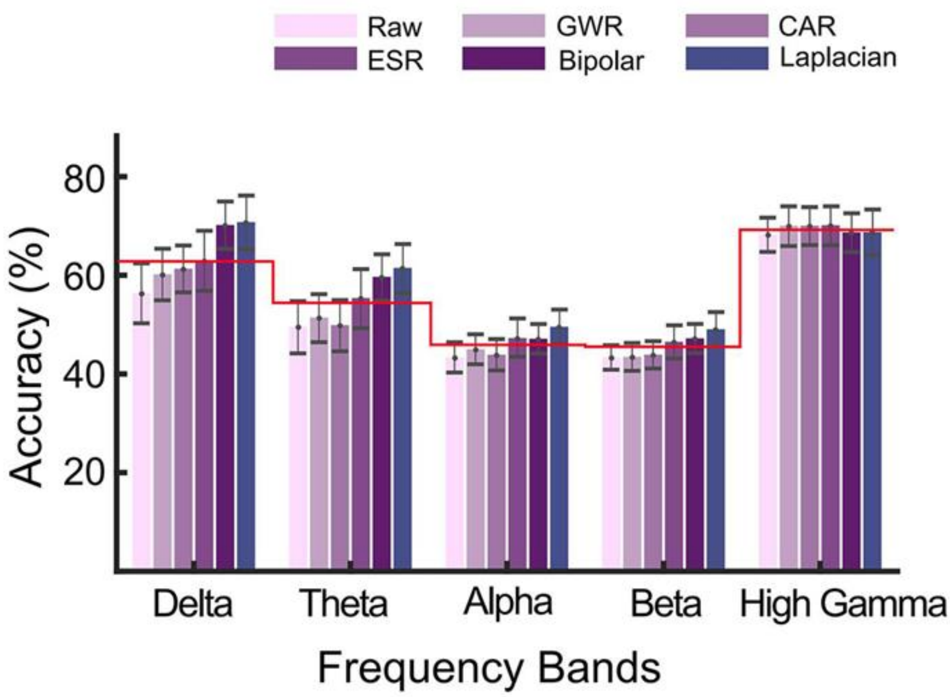

Figure 8: Single-frequency band-based decoding accuracy (SDA) of data cleaning methods for multiple sub-bands. Comparison of SDAs of different frequency bands. Bars and error bars represented mean accuracy and standard error calculated across all subjects, respectively. Redline represented average SDAs across all methods for different bands, respectively. Figure taken from [26].

3. Discussion of results

Overall, the way in which you re-reference your intracranial sEEG data will have a transformative effect on your signal. We gave an overview of three comparative studies that looked at sEEG data collected from epilepsy patients who were implanted with depth electrodes for the purpose of monitoring and observation. Each of these studies compared a range of re-referencing methods and their effect on the research objective in question. Two of these studies [14, 26] focused on motor task paradigms, with a focus on analysing data collected from cortical areas where such processes are typically observable. The dominant oscillatory rhythms in these regions are faster alpha and gamma frequencies [27], which might be locally produced by individual cortical regions as they respond to the environment. The third study focused on an associative memory paradigm, with a focus on analysing data collected from the medial temporal lobe, which is thought to be essential for memory processes [30]. This region is thought to be dominated by theta oscillations [15], a global rhythm that enables long range communication and memory processes [28, 31].

This overview therefore did not consider any similar comparative study over other domains or regions, such as the visual domain or prefrontal cortex. The explicit focus of this Chapter is time-frequency analyses and oscillations, so other important and understudied issues such as event-related potentials and related phasic phenomena were not considered. The predominant focus of the above analyses was on macro contact data for sEEG depth electrodes. A better understanding of how re-referencing effects sEEG data on grid and strip electrodes is still required. Another not well understood mode of data is that which is collected from micro-wires, an increasingly popular technique which can be used in tandem with the macro contacts and are capable of recording highly localised single unit and population activity (see also Chapters 49-53). In studying this issue, it might prove useful to understand how re-referencing might introduce phase-reversals into highly localised data and the likely ensuing effect that this has on the identified waveforms of neuronal spikes. Indeed, the issue of phase-reversals might be a prescient one for any analyses of oscillatory phase in general (which neighbour driven re-referencing methods might be more susceptible to) – which also requires further study.

In general, the re-referencing methods described here each work according to their own assumptions to emphasise either local or global effects, dependent on the type of data you are working with and the paradigm you are interested in. In general, re-referencing can be a good way to increase the signal to noise ratio of your data [26] – though it might prove insufficient to reduce sharp transient noise such as IEDs [25] – and might even contaminate previously unaffected signals in this respect. More localised methods (such as the neighbour-driven Laplacian or bipolar) might best enhance local effects, especially if spatial resolution is an important element in the design of the analyses [14]. However, simulations have shown (see Figure 4; [6]) that a data-driven method (such as independent component analyses or ICA), might be better suited to both isolate signal from local sources and discard activity from distant sources – especially in noisier conditions. Equally, more global methods (such as methods that make use of the average) might best enhance more global effects in the lower frequency range [25]. The data-driven ICA method seems to perform equally well at both isolating local signals [6], and enhancing global signals [25], depending on the hypothesis and parameters utilised – highlighting the versatility of data-driven methods to extract desirable components for further analyses.

In sum, more comparative work is required to build a full image of the transformative effect that re-referencing has on intracranial electrophysiological data that is collected from various neuronal sources and for a variety of hypothesis. Hopefully this Chapter will encourage researchers to consider re-referencing as an active pre-processing decision that needs careful consideration.

References

1. Hasselmo, M.E., *How we remember : brain mechanisms of episodic memory*. 2012, Cambridge, Mass. ; London: MIT Press. xii, 366 p.
2. Schomer, D.L. and F.H. Lopes da Silva, *Niedermeyer's electroencephalography : basic principles, clinical applications, and related fields*. Seventh edition. ed. 2018, New York, NY: Oxford University Press. xviii, 1239 pages.
3. Nunez, P.L., R. Srinivasan, and Oxford University Press., *Electric fields of the brain the neurophysics of EEG*. 2006, Oxford University Press,: Oxford. p. 1 online resource (xvi, 611 p.
4. Mercier, M.R., et al., *Evaluation of cortical local field potential diffusion in stereotactic electro-encephalography recordings: A glimpse on white matter signal.* Neuroimage, 2017. **147**: p. 219-232.
5. Lei, X. and K. Liao, *Understanding the Influences of EEG Reference: A Large-Scale Brain Network Perspective.* Front Neurosci, 2017. **11**: p. 205.
6. Michelmann, S., et al., *Data-driven re-referencing of intracranial EEG based on independent component analysis (ICA).* J Neurosci Methods, 2018. **307**: p. 125-137.
7. Staresina, B.P., et al., *Hippocampal pattern completion is linked to gamma power increases and alpha power decreases during recollection.* Elife, 2016. **5**.
8. Fell, J., et al., *Human memory formation is accompanied by rhinal-hippocampal coupling and decoupling.* Nat Neurosci, 2001. **4**(12): p. 1259-64.
9. Lachaux, J.P., D. Rudrauf, and P. Kahane, *Intracranial EEG and human brain mapping.* J Physiol Paris, 2003. **97**(4-6): p. 613-28.
10. Nunez, P.L., *REST: a good idea but not the gold standard.* Clin Neurophysiol, 2010. **121**(12): p. 2177-80.
11. Varela, F., et al., *The brainweb: phase synchronization and large-scale integration.* Nat Rev Neurosci, 2001. **2**(4): p. 229-39.
12. Treder, M.S., et al., *The hippocampus as the switchboard between perception and memory.* Proc Natl Acad Sci U S A, 2021. **118**(50).
13. Roux, F., et al., *Oscillations support short latency co-firing of neurons during human episodic memory formation.* biorXiv.org, 2022.

14. Li, G., et al., *Optimal referencing for stereo-electroencephalographic (SEEG) recordings.* Neuroimage, 2018. **183**: p. 327-335.
15. Griffiths, B.J., et al., *Directional coupling of slow and fast hippocampal gamma with neocortical alpha/beta oscillations in human episodic memory.* Proc Natl Acad Sci U S A, 2019. **116**(43): p. 21834-21842.
16. Comon, P., *Independent Component Analysis, a New Concept.* Signal Processing, 1994. **36**(3): p. 287-314.
17. Delorme, A. and S. Makeig, *EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis.* J Neurosci Methods, 2004. **134**(1): p. 9-21.
18. Whitmore, N.W. and S.C. Lin, *Unmasking local activity within local field potentials (LFPs) by removing distal electrical signals using independent component analysis.* Neuroimage, 2016. **132**: p. 79-92.
19. Nikulin, V.V., G. Nolte, and G. Curio, *A novel method for reliable and fast extraction of neuronal EEG/MEG oscillations on the basis of spatio-spectral decomposition.* Neuroimage, 2011. **55**(4): p. 1528-35.
20. Schaworonkow, N. and B. Voytek, *Enhancing oscillations in intracranial electrophysiological recordings with data-driven spatial filters.* PLoS Comput Biol, 2021. **17**(8): p. e1009298.
21. Liu, Y., et al., *Human Replay Spontaneously Reorganizes Experience.* Cell, 2019. **178**(3): p. 640-652 e14.
22. Kubanek, J., et al., *Decoding flexion of individual fingers using electrocorticographic signals in humans.* J Neural Eng, 2009. **6**(6): p. 066001.
23. McFarland, D.J., et al., *Spatial filter selection for EEG-based communication.* Electroencephalogr Clin Neurophysiol, 1997. **103**(3): p. 386-94.
24. Pfurtscheller, G., et al., *Mu rhythm (de)synchronization and EEG single-trial classification of different motor imagery tasks.* Neuroimage, 2006. **31**(1): p. 153-9.
25. Parish, G., et al., *The effect of re-referencing on low frequency power and spike-field coherence in intracranial EEG data.* in prep.
26. Liu, S., et al., *Investigating Data Cleaning Methods to Improve Performance of Brain-Computer Interfaces Based on Stereo-Electroencephalography.* Front Neurosci, 2021. **15**: p. 725384.
27. von Stein, A. and J. Sarnthein, *Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization.* Int J Psychophysiol, 2000. **38**(3): p. 301-13.
28. Buzsaki, G., *Theta oscillations in the hippocampus.* Neuron, 2002. **33**(3): p. 325-40.
29. Herweg, N.A., E.A. Solomon, and M.J. Kahana, *Theta Oscillations in Human Memory.* Trends Cogn Sci, 2020. **24**(3): p. 208-227.
30. Squire, L.R., C.E. Stark, and R.E. Clark, *The medial temporal lobe.* Annu Rev Neurosci, 2004. **27**: p. 279-306.
31. Solomon, E.A., et al., *Dynamic Theta Networks in the Human Medial Temporal Lobe Support Episodic Memory.* Curr Biol, 2019. **29**(7): p. 1100-1111 e4.