Investigating the limits of temporal clustering analysis for detecting epileptic activity in functional magnetic resonance imaging data

Yasha Borna Khatamian

Biomedical Engineering
McGill University, Montreal

August 2010

A thesis submitted to McGill University in partial fulfilment of the requirements of the degree of Master of Engineering

Abstract

Precise localization of epileptic activity is a necessity for those patients who may benefit from resective surgery. One common localization technique, EEG functional MRI (EEG-fMRI), localizes activity in an fMRI recording by finding blood oxygen level dependent (BOLD) signal correlates to epileptic events detected in a simultaneously recorded EEG. 2D temporal clustering analysis (2D-TCA) is a relatively new fMRI-based epileptic activity localization technique that breaks BOLD activity into components based on timing, finding epileptic activity without simultaneously recorded EEG. This study evaluated the ability of 2D-TCA to detect both simulated epileptic activity and activity detected in patients using EEG-fMRI. Although it was found that 2D-TCA could effectively detect epileptic activity with certain characteristics, it also detected activity not associated with epilepsy. As such, it was determined that 2D-TCA can only be used to validate epileptic activity localization by other means or to create hypotheses as to where activity may occur.

Résumé

Une localisation précise de l'activité épileptique est une nécessité pour les patients qui pourraient bénéficier d'une opération résective. L'EEG-IRM fonctionnelle (EEG-IRMf) est une nouvelle technique de localisation qui localise l'activité épileptique dans un enregistrement IRMf en trouvant un signal « blood oxygen level dependent » (BOLD) qui correspond à des événements épileptiques détectés simultanément par un enregistrement EEG. L'analyse temporelle groupée 2D (ATG-2D) est une technique de localisation de l'activité épileptique relativement nouvelle, qui est basée sur des données IRMf. Pour trouver l'activité épileptique, elle décompose l'activité BOLD en différentes composantes selon le moment où elles surviennent sans recourir à un enregistrement EEG simultané. Cette étude évalue la capacité de la technique ATG-2D à détecter une activité épileptique simulée ainsi que sa capacité à détecter une activité épileptique précédemment détectée chez des patients avec EEG-IRMf. Même nous avons montré que la technique ATG-2D pouvait détecter de façon efficace une activité épileptique ayant certaines caractéristiques, il a aussi été trouvé qu'elle détectait de l'activité non épileptique. Il a été déterminé que la technique ATG-2D pouvait seulement être utilisée pour valider une activité épileptique localisée par d'autres moyens ou pour formuler des hypothèses concernant l'endroit où l'activité pourrait survenir.

Acknowledgments

I would first like to thank Dr. Jean Gotman for his constant supervision of the work carried out for this thesis, his continuous support throughout the process, both financial and intellectual, and his ever-apparent commitment to the education of all members of his research group. Without his clear minded approach to research it is certain that obtaining my Master's degree would not have been nearly as enriching of a process as it was.

I would also like to thank all other members of epilepsy lab at the Montreal Neurological Institute for providing a warm and friendly environment within which research could be carried out. I would particularly like to thank Taha Gholipour and Firas Fahoum for being wonderful office mates and great resources for information on epilepsy, as well as Natasha Zazubovits for her continuously cheerful demeanour and help with data management; extra thanks to Firas for his analysis of patient data used in this thesis. Thanks also to Pierre LeVan whose rapid answers by e-mail to questions concerning computer code or analysis methods were greatly appreciated, Francesca Pittau for her help with questions concerning the EEG-fMRI data of patients, and Rina Zelmann for her help with anything engineering related, as well as her guidance on pursuing a Master's degree in Biomedical Engineering.

I would also like to thank Toula Papadopoulos, Lina Vuch, and Pina Sorrini for their always enthusiastic help with administrative issues.

Finally, I would like to thank all members of my family, including my sister in law Maëlle Turbide who helped translate my abstract into French, and particularly my parents, for their constant support, guidance, and tolerance during this and all other periods of my life.

This project was supported by a CGS M scholarship from the Natural Sciences and Engineering Research Council of Canada (NSERC) and by the Canadian Institutes of Health Research (CIHR) grant number MOP-38079.

Table of Contents

1 Introduction

Epilepsy is a neurological condition that affects approximately 0.7% of the population (Blume 2003) and can be described as a medical syndrome consisting of recurrent unprovoked seizures that can severely affect a patient's quality of life. While the majority of patients are able to achieve reasonable seizure control through use of medication, a large percentage remain refractory to such forms of treatment. In such cases resective surgery of the epileptogenic zone may be considered, a procedure that can lead to complete seizure control.

Electroencephalography (EEG) is the most common recording modality used to help in the diagnosis of epilepsy (Engel 1984), however in general it cannot localize an epileptogenic zone to the level of precision required by pre-surgical assessment. EEG-functional magnetic resonance imaging (EEG-fMRI) is technique that allows for precise localization of epileptic activity by looking for hemodynamic changes in the brain, recorded through fMRI, that correlate to epileptic events detected in the simultaneously recorded EEG of the patient. However, this technique is somewhat cumbersome as it requires the recording of EEG from within a running MR scanner. In addition, it is insensitive to epileptic activity restricted to deep brain structures as such activity would not be detected in the patient's EEG.

Temporal clustering analysis (TCA), and more specifically 2D-TCA, is a data based fMRI technique that has shown potential in being able to detect regions of the brain associated with epileptic activity with no dependency on EEG. This is done by clustering voxels based on the timing of activity detected in their time courses.

1.1 Project Objectives

Although it has been shown that 2D-TCA can detect epileptic activity in fMRI data without the need for an EEG recording, this technique has only been carried out in a handful of studies, most of which applied the technique to a relatively select number of patient scans. As such, the limits of its capabilities to detect various forms of epileptic activity in terms of activation size, amplitude, and frequency have not been evaluated. This is an important step in determining whether or not the general application of 2D-TCA to detect epileptic activity is actually worthwhile. The project presented in this thesis consisted of implementing an improved 2D-TCA algorithm and investigating its ability to detect epileptic activity of various activation sizes, amplitudes, and frequencies in both simulated fMRI data, as well as in fMRI recordings of epileptic patients. This was done to determine how effectively 2D-TCA is able to detect epileptic activity of various forms and to determine its true potential, in comparison to other techniques such as EEG-fMRI, for this application.

1.2 Thesis Organization

This thesis is organized as follows: Chapter 2 will consist of literature review that will cover background topics including epilepsy, EEG, fMRI, EEG-fMRI, and TCA. This will be followed by the actual investigation carried out for this thesis, chapter 3 covering methods employed, chapter 4 results obtained, and chapter 5 a discussion of those results, including possible next steps that could be taken. Chapter 6 will consist of a short conclusion summarizing what was investigated in the study and what was found.

2 Literature Review

2.1 Epilepsy

To understand what is meant by the term epilepsy, it is first important to have a clear definition of what a seizure is. The definition of a seizure, as accepted by the International League Against Epilepsy (ILAE), is “a transient occurrence of signs and/or symptoms due to abnormal excessive or synchronous neuronal activity in the brain” (Engel 2006). Approximately 5% of the general population will have a seizure at some point in their life (often in childhood), with common causes including sleep deprivation, hypoglycaemia, drug abuse, and head trauma, among others (Blume 2003). It is therefore important to distinguish between someone who has a seizure at some point in their life, and someone who is diagnosed with epilepsy. Epilepsy can be described as a medical syndrome consisting of recurrent unprovoked seizures and affects approximately 0.7% of the population (Blume 2003) (from this point on, unless otherwise stated, any use of the term seizure will refer to a seizure associated with a diagnosis of epilepsy). Someone with epilepsy is then thought of as having an epileptogenic abnormality that is intrinsic to the brain, exists between seizures, and causes seizures to occur independently of any acute state of the body (Engel 2006). This abnormality may cause random seizures in an individual for a short period of time, over many months, or even for an individual’s entire life and therefore can have a significantly detrimental effect on the physical and social well being of that individual.

Due to its general nature, epilepsy should be thought of as consisting of a wide set of syndromes. While these syndromes may each have their own symptoms, which are dependent on what anatomical structures in the brain are involved with the epilepsy (the epileptogenic zone), they are all considered as epilepsy in that they result from randomly timed abnormal electrical activity. Factors that

can help separate one epileptic syndrome from another include localization of the epileptogenic zone, seizure frequency, the age at onset of the epilepsy, and neurological history, among others (Blume 2003). Although epilepsy is considered to occur due to randomly timed neuronal activity, precipitants (including flashing lights, hyperventilation, sleep, sleep deprivation, and others) that can help induce a seizure often exist and can also be used to help identify the exact syndrome seen in a patient (Frucht et al 2000).

Epileptic syndromes and seizures can generally be considered as either "generalized" or "focal" (also referred to as "partial"). We now turn to defining what is meant by these two terms and will describe a few common syndromes associated with each.

2.1.1 Generalized Epilepsy

A generalized seizure refers to one that starts, or whose onset is, in widespread areas of both hemispheres of the brain. A syndrome is then referred to as generalized if the seizures associated with it are generalized. While all generalized seizures involve loss of consciousness in the individual, they can be subdivided based on the effect the seizure has on the individual's body. The following are a few examples of common forms of generalized seizures.

2.1.1.1 Tonic and/or Clonic Seizures

Also referred to as "grand mal" seizures, tonic and/or clonic seizures are likely the most common form of seizures associated with epilepsy by the general public. While each can consist of a seizure in itself (i.e. a tonic seizure, or a clonic seizure), a tonic-clonic seizure can be divided into the two phases, tonic and then clonic, and is often preceded by some sort of clinical symptom, called an aura, that can be very short, or last several hours. Patients who do not have an aura associated with their seizures are often thought of as having primary generalized seizures while those who do have an aura have secondary generalized seizures

(i.e. a seizure that becomes generalized after some initial focal activity; see section 2.1.2). The aura seen in a patient who has secondary generalized seizures is determined by where the initial seizure focus occurs; some common aurae include déjà vu, light-headedness, dizziness, unusual emotions, and altered vision and/or hearing. While in the tonic phase, which usually lasts only a few seconds, an individual will lose consciousness and then their skeletal muscles will tense, often causing the person to fall. In the clonic phase the individual's muscles will repeatedly contract and relax causing them to have convulsions (what is often mistakenly termed by the general public as a "seizure"). After the seizure has passed, the individual will often be confused for a period of time, not knowing what had happened and not remembering what they had been doing beforehand.

2.1.1.2 Absence Seizures

Also referred to as "petit mal" seizures, absence seizures often start in childhood or early adolescence (Blume 2003). An absence seizure consists of a sudden arrest of normal activity in an individual that lasts for approximately 5-20 s and then abruptly stops (Blume 2003). During this time, an outside observer would generally see the individual stare off into space as if without focus (in some cases some sort of unconscious movement such as blinking or a jerking arm might also be seen). After the seizure has passed, the individual will often continue as they were before hand, perhaps with a bit of initial confusion as to what had happened. Individuals who have absence seizures will typically have a few of these occurrences a day.

2.1.1.3 Other Forms of Generalized Seizures

Two other common forms of generalized seizures are myoclonic and atonic. A myoclonic seizure consists of the involuntary twitching of a muscle or a group of muscles; this differs from a clonic seizure in that a clonic seizure typically consists of much larger movements (e.g. a twitching lip compared to a full biting

motion). An atonic seizure consists of a complete, albeit brief, lapse in muscle tone; as the muscles completely relax this usually results in the person falling over.

2.1.2 Focal Epilepsy

A focal seizure refers to one whose onset occurs in part of one hemisphere of the brain (Engel 2006). A syndrome is then referred to as focal if the seizures associated with it are focal. Focal seizures can be divided into “simple”, in which the consciousness of the individual is unaffected, and “complex”, in which the consciousness of the individual is affected. It is also possible for a seizure to start focally, and then become generalized; such seizures are referred to as secondarily generalized. In fact, it is very unlikely that any form of seizure is strictly restricted to a focal region, independently from other areas; likewise, no seizure can be deemed as truly generalized (Engel 2006) in the sense that it may not involve absolutely the whole brain. We now turn to describing the most common form of focal epilepsy, temporal lobe epilepsy.

2.1.2.1 Temporal Lobe Epilepsy

The temporal lobe is the most common site for focal seizures (Blume 2003). Like tonic-clonic seizures, seizures associated with temporal lobe epilepsy (TLE) often begin in childhood or early adolescence (Blume 2003). In addition, this is the most common form of epilepsy seen in adults that does not respond to anti-epileptic drugs. For example, 89% of patients with mesial temporal sclerosis (one of the causes of temporal lobe epilepsy) are uncontrollable with medication (Semah et al 1998). Most seizures in individuals with temporal lobe epilepsy are complex, often consisting of an aura associated with neuronal activity occurring in the temporal lobe, for example an epigastric sensation, fear, or an olfactory phenomenon (Blume 2003).

2.1.3 Effect of Epilepsy on a Patient's Life

Epilepsy is one of the most common of the serious neurological disorders. Even in the less severe cases, due to the sporadic nature of epileptic seizures, a patient's independence can be severely restricted. Then depending on the actual severity of the condition many parts of the patient's life, whether it be their education, employment, social relationships, or sense of self-worth, can be seriously affected. Activities that people with epilepsy can often be forbidden to take part in include driving, working with heavy machinery, or any other task that requires continuous and uninterrupted awareness. For example, driving accidents are up to 1/3 greater among drivers with epilepsy (Blume 2003). In some cases however, exceptions to the rule can be made, for example with those patients who consistently have a long aura before their seizure occurs, or whose seizures only occur during sleep. In the case of driving, a risk assessment, that takes into account many factors such as seizure frequency and loss of awareness, must be completed to determine whether or not a licence should be issued or suspended (Blume 2003).

2.1.4 Treatment of Epilepsy

2.1.4.1 Anti-Epileptic Drugs

Most individuals with epilepsy find successful treatment with the use of anti-epileptic drugs (AEDs). With continuous use, AEDs can help reduce a patient's seizure frequency, in some cases completely stopping the seizures from occurring. This being said, AEDs can never cure a patient of their epilepsy. For example, if a specific AED is found to control a certain patient's epilepsy, but then they are taken off this medication, their seizures will reappear. In some cases multiple AEDs can be prescribed to a patient at the same time, but in most cases monotherapy is sufficient (Blume 2003). However, like any form of medication, side effects can result from the use of AEDs. 88% of patients using an AED report at least one side effect from usage of the drug (Baker et al 1997),

the most common side effect being fatigue (Blume 2003). As the use of an AED is typically for a prolonged time, often lifelong, any side effects can have a serious effect on the life of the patient.

When it is unclear what specific AED may be best suited to control a patient's epilepsy, a set of probable candidates would be tested. Typically, a given AED would be tested by gradually increasing the dosage given to the patient until it works, or until any of its side effects become too intense to warrant further usage (Noachtar & Borggraefe 2009). One study showed that 47% of patients will achieve seizure control with the first AED tested, 14% with a second or third tested AED, and 3% with the use of two AEDs at the same time (Birbeck et al 2002). To deem a patient's epilepsy as medically intractable at least two or three AEDs must be shown to be ineffective when applied in monotherapy (Bourgeois 2001). Approximately 30% of patients with epilepsy are found to be intractable to AEDs (Kwan & Brodie 2000). In these situations other methods, such as surgery, may be considered depending on how suitable a given patient's case may be.

2.1.4.2 Surgical Treatment

The ultimate goal for surgical treatment is to achieve complete control of a patient's epileptic seizures while minimizing any detrimental effect to the patient's normal function. In some cases, where the patient no longer has seizures after the surgery, the treatment can be considered as curative, while in others the seizure frequency is greatly reduced post-operatively, and may then be controlled through the use of medication (Berg et al 2007).

The first step in even considering surgery for a given patient is to confidently establish the medical intractability of the patient's epilepsy. After this crucial step has been taken, the candidacy of the patient's case for surgery must be determined. Criteria that must be met include: disabling seizures, motivated

patient, and a high probability that better seizure control will improve the patient's quality of life (Noachtar & Borggraefe 2009). In addition, as the results of surgery generally depend on how well the epileptogenic zone can be localized, and how effectively it can be removed without damaging functionally essential brain structures, determining the candidacy of a patient mostly relies on the result of a rigorous pre-surgical evaluation consisting of many tests whose goal is to confidently localize the epileptogenic zone. These tests typically include: long-term video-EEG monitoring, often over days, to record clinical and electrical manifestations of the epilepsy; neuroimaging, for example anatomical magnetic resonance imaging (MRI); and, neuropsychological tests. In some cases, when deemed necessary, invasive EEG may also be used if more precise localization is required, however, for 80-90% of patients who are being considered for surgery, non-invasive methods are sufficient to determine whether they are suitable or not (Noachtar & Borggraefe 2009). In the case of lesional epilepsies, an anatomical MRI can effectively indicate the lesion, which is often where the epileptogenic zone exists. On the other hand, the localization of a non-lesional epileptogenic zone must rely solely on functional tests.

As mentioned before, TLE is the most common form of epilepsy seen in adults that does not respond to anti-epileptic drugs. Fortunately, surgery as treatment for TLE often provides effective and consistent results, even potentially reducing health costs (Wiebe et al 2001). One study found that 70% of TLE patients who undergo surgery become seizure free post-operatively, while 20% see a significant reduction in seizure frequency (Engel 1993). This high success rate is largely due to the fact that the limits of resection for TLE are often easier to define as compared to other forms of epilepsy for which the boundaries of the epileptogenic zone are more variable. In some cases, the goal of surgery may be to simply reduce the effects of a seizure on the state of a patient. For example, a callosotomy, which consists of disconnecting certain pathways between the two

cerebral hemispheres, may be carried out to prevent a patient's seizures from generalizing to their entire brain, which would cause them to lose consciousness.

2.2 EEG

EEG is the most common recording modality used to help in the diagnosis of epilepsy (Engel 1984). Though the diagnosis of epilepsy is typically made based on clinical observations, an EEG recording can help confirm that diagnosis and may then be able to characterize the specific syndrome. For example, various EEG patterns can be correlated to specific epilepsy syndromes, and therefore an EEG recording of a patient is very useful for prognosis and in determining the specific form of therapy that should be taken for a given patient (Noachtar & Remi 2009). Some of the benefits of EEG compared to other recording modalities, and reasons that it is so widely used for epilepsy, include its relative low cost, high temporal resolution, ability to directly record neuronal electrical activity, ability to be paired with simultaneous video recording of seizures, and the fact that patients are relatively free to move and will therefore have no issues with claustrophobia (particularly useful for prolonged recordings).

2.2.1 Basis of the EEG Signal

Figure 2.2-1 shows the main components of a neuron, the basic building block cell of the brain. It consists of dendrites, a cell body, an axon, and axon terminals. Neurons pass information to one another through use of electrical signals. For example, if the difference between the electrical potential inside and outside the cell body exceeds a certain voltage threshold (around -30 mV, the cell resting potential is typically -70 mV) the cell will become activated and an action potential, which essentially consists of a discrete electrical impulse, will be sent down the axon from the cell body. This impulse will then reach the axon terminals of the activated cell, which will connect to the dendrites of other neurons, each through a synapse (the point of connection between the terminal of the first neuron, deemed the pre-synaptic neuron, and dendrite of the second neuron, deemed the post-synaptic neuron). When the action potential reaches a synapse, the axon terminal will release neurotransmitters to the open area

between the terminal and dendrite, causing an electrical signal to be induced in the dendrite of the post-synaptic neuron. If the combination of electrical potentials in the dendrites of the post-synaptic neuron exceeds threshold, its cell body will send an action potential and the signal will continue down its axon to other neurons. It should be noted that one neuron can receive electrical signals from, and send electrical signals to, many neurons.

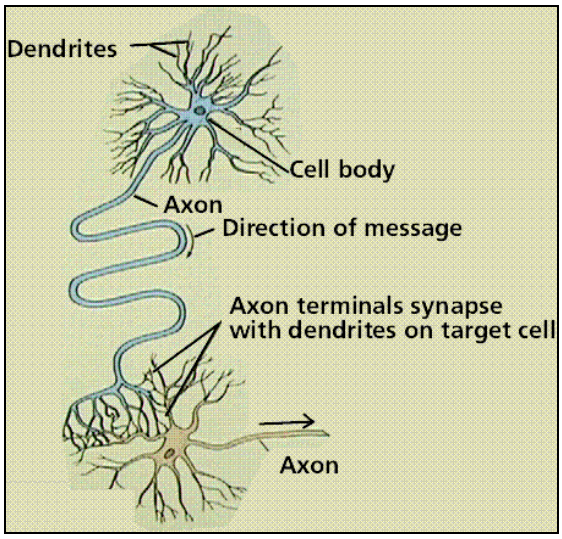

Figure 2.2-1 Structure of a neuron and the direction of nerve message transmission. Reproduced from (Purves 2004).

While one might think that EEG, which most commonly consists of electrodes placed on the scalp, records signals associated with the action potential, it is actually insensitive to these signals as they only cause local currents that do not penetrate far past the space surrounding the activated cell; these currents are also not likely to summate because of their short duration. EEG actually records electric fields associated with the post-synaptic potentials seen in neurons. However, this electrical field must be strong enough to reach the electrodes. As such, for EEG to be able to detect an electric field two criteria must be met: 1) the post-synaptic terminals of a large number of neurons must be activated at the same time, and 2) these synchronous neurons must be spatially oriented in a similar direction so that there is little interference between their potentials. If

these two criteria are met the field of the sum of these neurons will be strong enough to reach the electrodes and the signal will be able to be picked up by the EEG.

Unfortunately only post-synaptic potentials from cortical neurons are detectable, and activity deep within the brain is generally not picked up. This is because electrical fields fall off with the square of distance, and because those signals originating deep in the brain are deteriorated by having to pass through a good portion of the brain, in addition to the meninges, the skull, and the scalp. These structures, especially the skull, also diffuse the signal causing deterioration in the overall spatial resolution achievable with EEG.

2.2.2 EEG Recording and Analysis

As mentioned, EEG typically consists of recording signals by placing electrodes on the scalp. These electrodes are placed using a conductive paste after the points of attachment on the scalp have been cleaned to remove dirt and dead skin that might impede the signal. Electrodes are typically placed in a very organized set of locations that adhere to the 10-20 system, as shown in Figure 2.2-2 (it should be noted that the system allows for more electrodes to be placed than are shown in Figure 2.2-2). Adhering to this system ensures that the location and naming of the electrodes are consistent from one EEG lab/clinic to another. Each of these electrodes is then connected to a differential amplifier to which a common reference electrode is also connected. The difference signals are then passed through an anti-aliasing filter, and recorded using an analog to digital converter. A typical EEG consists of signals on the order of 30 μV in amplitude (Aurlien et al 2004).

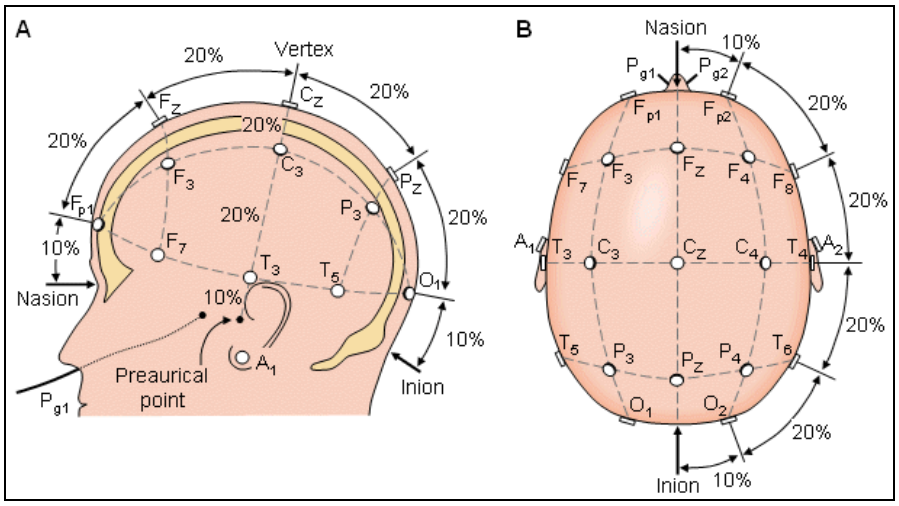

Figure 2.2-2 The 10-20 system of EEG electrode placement as viewed from (A) the left, and (B) the top. Reproduced from(Malmivuo & Plonsey 1995).

As the voltage signal associated with each electrode is actually a difference signal, a single EEG recording is often viewed in a variety of ways, each consisting of plotting out a specific set of differences between the various electrodes. In the case of epilepsy, analysis of an EEG recording typically consists of an electroencephalographer visually inspecting the recording, after it has been recorded, to detect any activity of interest, i.e. neuronal activity associated with epilepsy. To aid in increasing the chance that such activity is present when the actual recording of a patient is taking place, activation techniques, such as hyperventilation, photic stimulation (i.e. flashing lights), sleep, and sleep deprivation, may be used to induce the abnormal epileptic activity.

2.2.3 Interictal Epileptic Discharges

An epileptic patient can be considered as existing in one of two states, ictal and interictal. An ictal state refers to the period of time when the patient is actually having a seizure, while an interictal state refers to the period of time in between seizures. While an EEG is able to record neuronal activity associated with a seizure (Figure 2.2-3A), it can also pick up abnormal interictal activity (Figure 2.2-3B), termed interictal epileptic discharges (IEDs), that are typically only seen in epilepsy patients. Although no hard definition for an IED exists, and even

electroencephalographers can disagree on what is and is not an IED (Noachtar et al 1999), this discharge usually consists of a short burst or spike of activity from the same area as the ictal (seizure) activity is believe to originate. It should be noted that unlike ictal activity, IEDs generally do not produce any clinical manifestations.

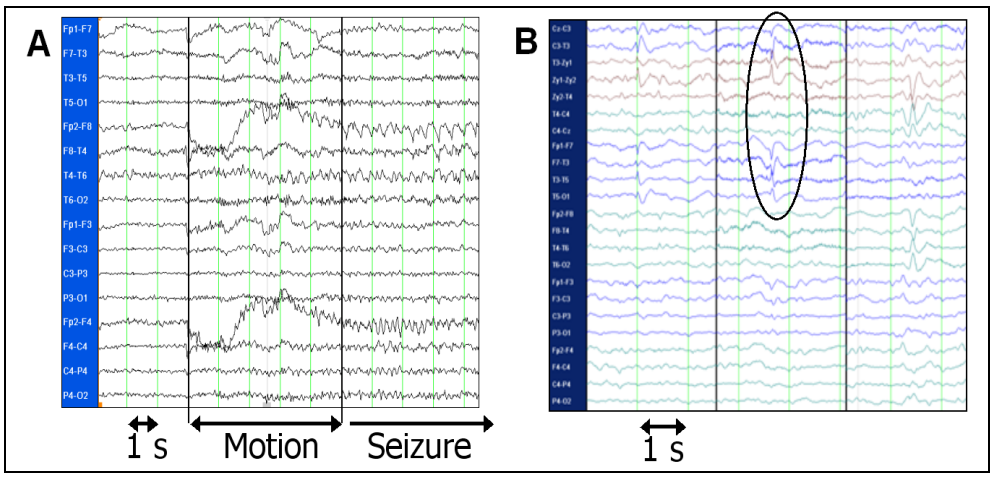

Figure 2.2-3 Example EEG recordings of (A) a seizure, and (B) an IED.

In addition to helping identify the specific syndrome of epilepsy in a given patient, the detection of IEDs is a very good indicator of whether an individual's seizures are epileptic or if they arise from some other non-epileptic condition (Noachtar & Remi 2009). Only about 0.5% of the general population show IEDs in their EEG (Gregory et al 1993) compared to about 98% of epilepsy patients (Marsan & Zivin 1970). Long-term recordings as well as repeated EEGs, both done to increase the likelihood of recording an IED, are required for any conclusive diagnostics as it is not uncommon for an epilepsy patient to have a normal EEG recording (Noachtar & Remi 2009) or for non-epileptic EEG patterns to be incorrectly interpreted as IEDs. In addition, the same techniques implemented to induce ictal activity during an EEG recording, as described in section 2.2.2, also increase the likelihood of recording an IED (Noachtar & Remi 2009).

2.2.4 Localizing Epileptic Activity Using EEG

Precise localization of ictal and interictal activity, often through the use of an EEG recording, are crucial steps in defining the epileptogenic zone in a patient, especially for those who are candidates for resective surgery (Noachtar et al 2003). Ictal recordings, especially when acquired simultaneously with video recordings, provide the best results (Noachtar & Remi 2009). When a regular scalp EEG does not provide precise enough results the use of invasive EEG may be considered, however, this depends on the given case.

2.2.4.1 Video-EEG Recording

While a routine EEG recording lasts around 30 minutes, to make sure that ictal activity is recorded, the EEG of an epilepsy patient may be continuously acquired over a considerably longer time, up to days. When possible, for example while the patient is in bed in their hospital room, video that is synchronized with the EEG recording is also recorded. These recordings can help characterize the syndrome of epilepsy within the patient so that appropriate treatment can be carried out. In addition, as mentioned in section 2.1.4.2, video-EEG recording is often one of the modalities used to help spatially localize the epileptogenic zone as part of a pre-surgery evaluation. The precision of localization required by a pre-surgery evaluation is not always achievable solely through use of a video and scalp EEG recording due to the EEG's low spatial resolution.

2.2.4.2 Inverse Methods

As mentioned, visual interpretation of the scalp EEG by an electroencephalographer provides relatively imprecise localization, often only being able to confirm within which lobe the activity is taking place. The goal of using inverse methods is to precisely localize, using mathematical and physical models, the source of the epileptic activity given what is seen in the electrodes of an EEG recording. Investigations into the ability to localize sources via inverse methods have only become more common in the last decade, largely due to

advancements in computer processing and storage capabilities (Plummer et al 2008). Applying inverse methods generally consists of the development and application of two models, forward and inverse.

A forward model aims to solve the forward problem which consists of simulating the activity seen in the electrodes of an EEG recording given the source of the activity. This is solved by specifying a set of conditions for the head that often consists of defining the different compartments, surfaces, and conductivities seen within it. Forward models (Kybic et al 2006; Mosher et al 1999) can be as simple as a single spherical shell, or as complex as a four compartment model (i.e. brain, cerebral spinal fluid (CSF), skull, and scalp) where boundaries are defined by segmentation of the subject's anatomical MRI scan. For a specific source of activity the forward model will simulate a specific distribution of activity in the electrodes, and as such can be thought as always providing a unique solution (Wilson & Bayley 1950).

An inverse model aims to solve the inverse problem which consists of localizing the source of activity given what is seen in the electrodes of an EEG recording. Methods to create the inverse model can usually be divided into two groups: 1) equivalent current dipole methods, which assume that EEG activity is generated by a few dipolar sources, and 2) distributed source methods, which assume that EEG activity is generated by a large number of dipolar sources distributed within the brain or on the cortical surface. While equivalent dipole methods are most useful in the localization of IEDs (Merlet & Gotman 1999), distributed source methods are most useful in localizing spatially extended sources (Grova et al 2006). Unfortunately, the inverse problem is in fact ill-posed as, theoretically, an infinite number of source configurations can produce the same electrical fields at the surface (Helmholtz 1853), meaning the inverse problem has no unique solution. As such, additional constraints/assumptions, both anatomical, for

example restricting possible source activity only to the cortical surface, and mathematical, for example finding the minimum energy solution, must be used to obtain a unique solution (Grova et al 2006).

2.2.4.3 Invasive EEG

In most cases, non-invasive recording methods provide precise enough results in the pre-surgical evaluation of a patient (Noachtar et al 2003), especially when imaging studies (i.e. MRI, positron emission tomography (PET)) find results congruent with what was expected based on the EEG and clinical observations (Noachtar & Remi 2009). Invasive recordings are therefore only used when either the non-invasive techniques do not provide precise enough localization (but provide enough information to have an idea as to where invasive electrodes should be placed), or if the epileptogenic zone is found to be very close to functionally essential cortex (Noachtar et al 2003).

Invasive EEG consists of using one of three electrode types, epidural, subdural, or depth, each probing deeper into the brain than the previous (see Figure 2.2-4). Epidural electrodes are placed, often through boreholes, just underneath the skull on top of the dura mater. Subdural electrodes often consist of a grid or strip of electrodes that are placed just under the dura mater on top of the cortex and therefore require a craniotomy. Depth electrodes are able to penetrate into the brain to record from deep structures, in addition to superficial ones, however trajectories for placement must be devised so that important brain structures and blood vessels are not damaged. Subdural and depth electrodes both provide a better signal to noise ratio (SNR) than epidural electrodes (Noachtar & Remi 2009), and all of these invasive techniques provide a better SNR than scalp EEG (on the order of mV rather than μV). For this reason, invasive EEG is able to detect activity that might be undetected by a scalp EEG (Noachtar & Remi 2009). However, in the case of subdural and depth electrodes it is crucial

that they are placed within a few mm of the epileptogenic zone, anything further and the recording will only show the spread of the epileptic activity and not the origin. In addition, like anything else that is invasive, the use of invasive EEG is less than ideal as 1-4% of patients have complications (i.e. hemorrhages, infection), most often temporary, that arise from the required surgery (Noachtar & Remi 2009).

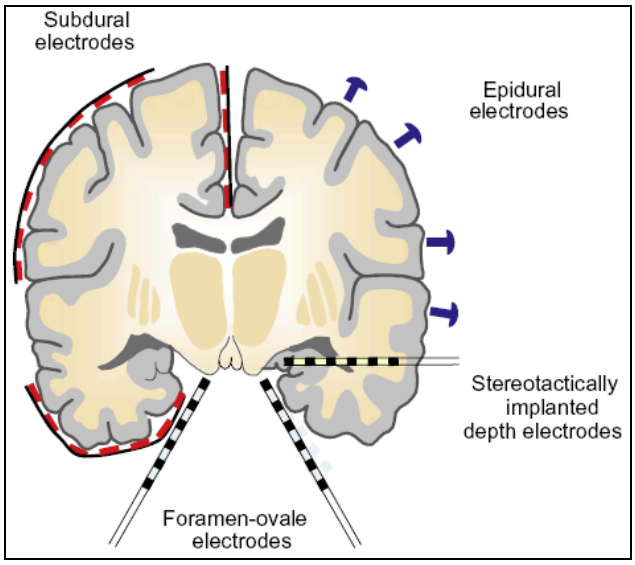

Figure 2.2-4 Placement of various invasive EEG electrodes. Reproduced from (Noachtar & Remi 2009).

2.3 Functional Magnetic Resonance Imaging

Magnetic resonance imaging (MRI) is a non-invasive technique that is able to acquire detailed 3D images of the brain and other body parts. It first evolved in the 1970s and has seen considerable growth since then, particularly as a diagnostic tool (Prasad & Storey 2008). While it is most often used to obtain a static image of the brain, which allows for analysis of brain anatomy, with the development of functional MRI (fMRI) in the 1990s it can now be used to record how the hemodynamics of the brain change with time. This section will review the physical basis of how an MR signal is actually acquired in addition to the nature of the signal recorded by fMRI and how it is typically analyzed.

2.3.1 Physical Basis and Acquisition of the MR Signal

MRI functions by exploiting the property of nuclear magnetic resonance seen in some atomic nuclei (i.e. a nucleus exposed to a strong static magnetic field will absorb and re-emit electromagnetic waves that resonate with it). As the resonant frequency of a nucleus is a precise measurement of its local magnetic field it is a good indicator of its molecular environment. This resonance seen in some nuclei arises from the fact that they possess a spin, or magnetic moment, which creates a small magnetic field around them, much like that seen in a bar magnet. Therefore, as shown in Figure 2.3-1, when a static external magnetic field is applied, the nucleus will begin to precess around the axis of the applied field, much like a wobbling spinning top under the influence of gravity. This precession occurs because the magnetic moment of the nucleus does not immediately align to the external magnetic field due to its spin, which continues to influence the direction of its magnetic moment.

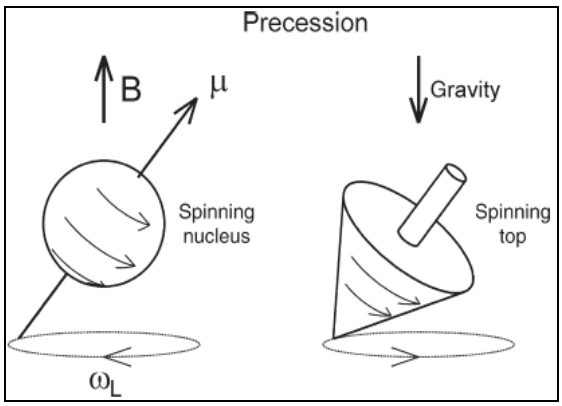

Figure 2.3-1 Precession seen in a nucleus with spin when an external field is applied, as well as precession seen in a spinning top under the influence of gravity. Reproduced from (Prasad & Storey 2008).

The frequency of precession seen in a nucleus is known as the Larmor frequency and is easily calculated as $\omega_L = \gamma B_0$ where B_0 is the strength of the externally applied field, and γ is the gyromagnetic ratio, a property specific to the nucleus type (e.g. γ = 42.58 MHz/T for hydrogen (Prasad & Storey 2008), the atom most commonly exploited in MRI). In a volume of tissue with many precessing nuclei, there will be a net magnetization consisting of a longitudinal component (along the axis of the applied field) and a transverse component (perpendicular to the applied field). When only the constant field B_0 is applied, there will be a slight net longitudinal component in the direction of the applied field, however no net transverse component will exist as all the nuclei, precessing at a frequency of ω_L, will be in random phase. By applying an external field in the transverse direction (perpendicular to the direction of B_0) that oscillates at the Larmor frequency, all nuclei will come into phase and a net transverse field, which also oscillates at the Larmor frequency, can be detected; this process is referred to as radio frequency (RF) excitation.

After the external RF pulse has finished the excitation process, the net transverse oscillating field in the tissue will slowly decay as the magnetic moments of the nuclei again start to fall out of phase due to random processes; this is referred to as transverse relaxation. In addition, the increase in the

longitudinal component caused by the RF pulse will subside due to energy loss to the surroundings; this is referred to as longitudinal relaxation. The timescale on which longitudinal relaxation occurs is captured in T_1, a time constant value that differs from tissue to tissue. The timescale on which the transverse relaxation occurs (which is considerably shorter than T_1) is captured in the value T_2. T_2 relaxation occurs as a result of dephasing between the nuclei. This dephasing is caused by inhomogeneities in the local field due to the interactions between neighbouring nuclei and molecules. While the value of T_2 describes the relaxation that occurs due to local interactions, inhomogeneities in the externally applied field will also cause dephasing to occur and will speed up the relaxation process. The effect of both dephasing mechanisms is captured in the time constant value T_2*. Figure 2.3-2 shows the relaxation in the net magnetic field according to T_2*, as well as the associated signal recorded perpendicular to B_0 by the MR scanner.

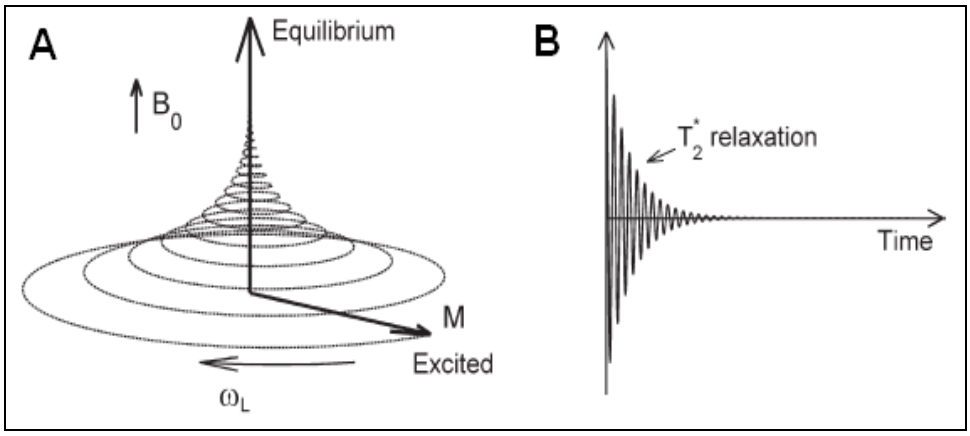

Figure 2.3-2 (A) Relaxation of net magnetic field after excitation and (B) associated recorded signal which decays according to T_2*. Adapted from (Prasad & Storey 2008).

To acquire data from a single slice a gradient field, that causes B_0 to vary along one dimension of space, can be applied, insuring that only those nuclei within the slice will be excited as only they will have a resonant frequency equal to that of the RF pulse to be applied. Likewise, additional gradients can be applied along the two dimensions of the slice to spatially encode that data so that it can be

recorded in k-space, i.e. the spatial frequency domain. The real image of the acquired slice can then be obtained by simply taking the inverse Fourier transform of the k-space image.

As the values of T_1, T_2, and T_2* vary from tissue to tissue, by recording a function that is an appropriate combination of these parameters MRI acquisitions can be tailored so that they are sensitive to specific tissues or molecules. For example, while T_1 and T_2-weighted images (i.e. recordings that are functions of T_1 and/or T_2 decay) show good contrasts between grey matter, white matter, and CSF, T_2*-weighted images are sensitive to areas of magnetic susceptibility making them particularly useful in detecting changes in the hemodynamics of the brain, i.e. what is detected in fMRI.

2.3.2 Vascular Response to Neuronal Activity

One of the major changes caused by neuronal activity is that seen in the local cerebral hemodynamics (Ogawa et al 1993), i.e. the cerebral blood flow (CBF), cerebral blood volume (CBV), and blood-oxygen concentration. These changes are a result of the brain's vascular response to neuronal activity. When a neuron becomes active it increases its energy consumption. This increase in energy consumption leads to an increase in oxygen extraction from the local blood source which in part increases the deoxyhemoglobin (dHb) concentration in the venous blood leaving the area of activity. This increase in dHb concentration then causes a vascular response to occur which consists of the dilation of the arterioles that feed the activated region, in part increasing the blood flow to the active area. With this increase in blood flow there is seen a decrease in the dHb concentration in the draining capillaries and venules (i.e. due to the large amount of oxyhemoglobin (Hb) flowing through). It takes a few seconds for these events to transpire.

While one may think that the hemodynamics will return to resting state levels once the neuronal activity has finished, this is not true. In actuality, the peak of the hemodynamic response will be followed by a period of time called the post-stimulus undershoot, during which the concentration of dHb in the venous blood is actually slightly above what was seen in the resting state. This undershoot typically starts between 10-20 s after the initial stimulus (Fransson et al 1999), and can last up to 60 s (Bandettini et al 1997). While the cause of this undershoot has not been confirmed, there are a few potential theories (Chen & Pike 2009). One, referred to as the biomechanical model, attributes the increase in dHb to a temporary mismatch between the CBF and CBV responses during the post-stimulus undershoot period (Mandeville et al 1999). While the arterioles recover quickly, causing the CBF to do the same, the venules recover more slowly, causing the venous blood volume to remain above resting state levels for a period of time, in part creating the apparent increase in dHb associated with the undershoot. Another theory attributes the undershoot to a decoupling between the CBF and level of oxygen extraction, i.e. while the CBV and CBF return to normal levels relatively quickly, oxygen extraction levels may remain slightly elevated for a period of time in part increasing the amount of dHb in the blood (Frahm et al 1996). A third, more recent theory, attributes the post-stimulus undershoot to a post-stimulus undershoot in the CBF (Behzadi & Liu 2005).

2.3.3 The Blood Oxygen Level Dependent Signal

As Hb is diamagnetic, while dHb is paramagnetic, the overall magnetic property of blood depends on the level of its oxygenation and therefore on physiology (Pauling & Coryell 1936). Due to its paramagnetic nature, dHb in the capillaries and venules distorts the local magnetic field causing hydrogen atoms in water to dephase and therefore the signal strength in that area to decrease. As such, the recorded blood oxygen level dependent (BOLD) signal is inversely affected by the

concentration of dHb, which as indicated in section 2.3.2, is dependent on the local oxygen extraction, CBF, and CBV. BOLD fMRI is by far the most commonly used fMRI technique and is usually acquired with a gradient-echo scan sequence which is sensitive to the susceptibility changes caused by changes in the concentration of dHb (Ogawa et al 1990).

2.3.4 The Hemodynamic Response Function

Under many conditions, the BOLD response to neuronal activity has been found to be approximately linear (Boynton et al 1996). The goal of the development of a hemodynamic response function (HRF) is to effectively model the BOLD signal response to an impulse of neuronal activity. This would then allow one to find, by simple convolution, the expected BOLD signal that would result from any time course of neuronal activity. The first model of the HRF was developed using a Poisson distribution function and was found to adequately describe responses in the primary visual cortex (Friston et al 1994). More recent models consist of a gamma function (i.e. a delayed peak) (Lange & Zeger 1997) or a difference of two gamma functions, like the commonly used Glover HRF (Glover 1999), which allows for the modeling of the post-stimulus undershoot as described in section 2.3.2. Figure 2.3-3 shows a plot of the Glover HRF, along with the stages associated with different parts of the vascular response. Although not modeled in the Glover HRF as it is essentially negligible, stage 1 consists of an initial dip in the BOLD signal below the baseline value of 0; this dip in the BOLD signal is a result of the increase in the dHb concentration due to the increase in oxygen extraction by the activated neurons. The vascular response (stage 2) to this initial dip is then seen as an increase in the BOLD signal (decrease in dHb concentration) caused by the increase in CBF and CBV. After the vascular response has subsided, the post-stimulus undershoot (stage 3) occurs until the signal returns to baseline after about 20 s.

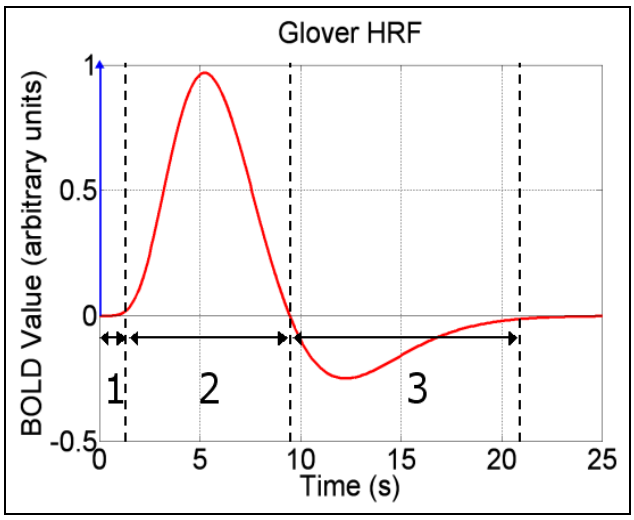

Figure 2.3-3 Plot of Glover HRF with various stages indicated: 1) initial dip (not modeled), 2) vascular response, 3) post-stimulus undershoot.

Although in practice it is often assumed that the HRF is identical for all individuals and in all areas of the brain, this is not necessarily true. In one study it was shown that the response in the sensory motor cortex varied across individuals (Aguirre et al 1998), while another study showed that in a single individual the impulse response derived from activity in the visual cortex differed from that derived from the auditory cortex (Robson et al 1998). Different HRFs are still relatively similar with the delay or peak time being the parameter that often makes the biggest difference in analysis.

2.3.5 GLM and Statistical Analysis of fMRI Data

Due to the nature of MRI, the BOLD signal can be recorded throughout the brain. However, as fMRI does not obtain a static image but rather a sequence of images that show changes in the BOLD signal over time, each volume must be acquired relatively quickly, usually accomplished by sacrificing the spatial resolution of the final scan. A typical fMRI scan contains around 25 slices acquired over a one to two-second window, each slice consisting of a 64x64 pixel image (total of 102 400 voxels, about 10 000 of which are in the brain) with voxel sizes of about 4x4x4 mm^3; each voxel then has its own BOLD signal time course sampled at around 0.5 Hz.

The change seen in an MR signal due to the BOLD effect is actually relatively small. For example, in tasks involving the primary or sensory motor areas the BOLD signal increases by less than 5% as compared to baseline, while memory tasks see an increase of less than 1% (Huettel et al 2004). As such, to get any reliable results statistical analysis methods must be employed. While early fMRI analysis techniques consisted of simply subtracting a resting state recording from an active state recording to determine changes associated with the active state, the most common analysis technique in use today is by far the general linear model (GLM).

In short, the aim of the GLM is to see how well the expected BOLD response to a given task describes the time course of a given voxel; those voxels who are well described are then deemed as activated by the task in question. As shown in Figure 2.3-4, the expected response is typically created by convolving an HRF with an input function that describes the states of activity associated with the task (e.g. 0 indicating a resting state and 1 indicating an active state, for example tapping a finger). The resulting time course then describes how the BOLD signal associated with the given task is expected to vary.

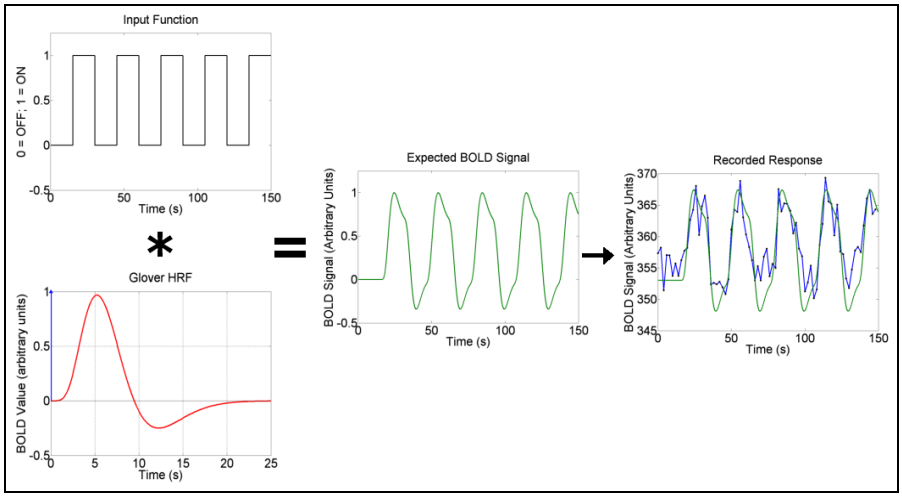

Figure 2.3-4 Convolution of input function and HRF to create BOLD response expected to result from given input function task. This BOLD response is then linearly fit to a given voxel time course.

One of the very powerful characteristics of the GLM is that a number of different regressors, each describing a different component of the fMRI data, can be easily added to or removed from the model. The final model consists of describing each voxel's time course by a linear combination of all appropriate regressors as shown in the equation:

$$\boldsymbol{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{e} = \boldsymbol{x}_1\beta_1 + \ldots + \boldsymbol{x}_n\beta_n + \boldsymbol{e}$$

where $\boldsymbol{y}$ is a column vector containing the time course of a given voxel, $\mathbf{X}$ is the design matrix containing all appropriate regressors as column vectors (i.e. $\boldsymbol{x}_1$, …, $\boldsymbol{x}_n$), $\boldsymbol{\beta}$ is a column vector containing scale factors for each regressor (i.e. β_1, …, β_n), and $\boldsymbol{e}$ is a column vector consisting of the noise not accounted for by the model. For example, in the very simple case of Figure 2.3-4, the voxel time course may be modeled, as shown in Figure 2.3-5, by a linear combination of the expected response $\boldsymbol{x}_1$, a constant valued regressor $\boldsymbol{x}_2$, and the error $\boldsymbol{e}$. To find the most appropriate model using a given set of regressors, the β values must be estimated such that the amount of variance in the voxel's time course described by the error, $\boldsymbol{e}$, is minimized. This can be achieved through ordinary least squares estimation, i.e. $\boldsymbol{\beta} = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\boldsymbol{y}$. Each voxel then has its own β value for each regressor.

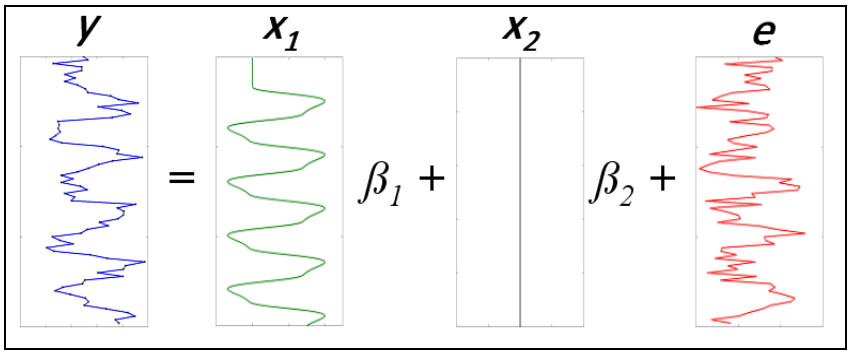

Figure 2.3-5 Modelling a voxel time course as the linear combination of an expected response, a constant signal, and error.

As previously mentioned, a voxel is deemed as activated by a given task when the expected BOLD signal response to that task describes its time course well

enough. This is determined through a t-test, which determines the probability that the voxel's time course is actually from the null hypothesis (i.e. that it is not activated) while taking into account its intrinsic variation and the model being used (i.e. the calculated β values). In simple terms, a t-value can be thought of as a ratio of the effect in the voxel's time course explained by the given model to the variance in the data that is unexplained by the model. Therefore, the higher the t-value is, the less likely that the voxel is not active.

A t-map is obtained by calculating the t-value associated with every voxel time course. A t-value threshold must then be implemented to differentiate between which voxels show statistically significant activation and which do not. In conventional statistical analysis, a t-value that corresponds to a p value of 0.05 is chosen (i.e. the probability that a voxel would have a t-value above threshold and would be deemed as activated purely by chance is 0.05 or less). However, considering that there is approximately 10000 voxels in the brain, a multiple comparisons problem arises as this could result in around 500 false positives purely by chance. Many techniques exist which find a more appropriate t-value threshold by taking into account various data parameters. For example, Bonferroni correction, which assumes that voxel time courses are mutually independent, addresses the issue by dividing the initial p value (0.05) by the number of voxels (10 000) to obtain a new p value of 0.000005. This technique is often thought to be too conservative due to the fact that the signals at each voxel are not actually mutually independent but rather have some spatial correlation caused by intrinsic properties and smoothing applied in pre-processing (see section 2.3.6.4). Random field theory is one technique that takes this smoothness into account by calculating an appropriate t-value based on the number of "resolution elements" rather than the number of voxels (Worsley et al 1996). In another technique, rather than controlling the number of expected false positives, the false discovery rate (ratio of false positives to total number of

positives found) is controlled to find an appropriate t-value threshold (Genovese et al 2002). In general, it is not uncommon to choose a p value of 0.001, which corresponds to a t-value threshold of 3.1. The anatomical region associated with the given task can then be easily visualized by overlaying the t-map on an anatomical scan of the same subject and applying the t-value threshold (as seen in Figure 2.3-6).

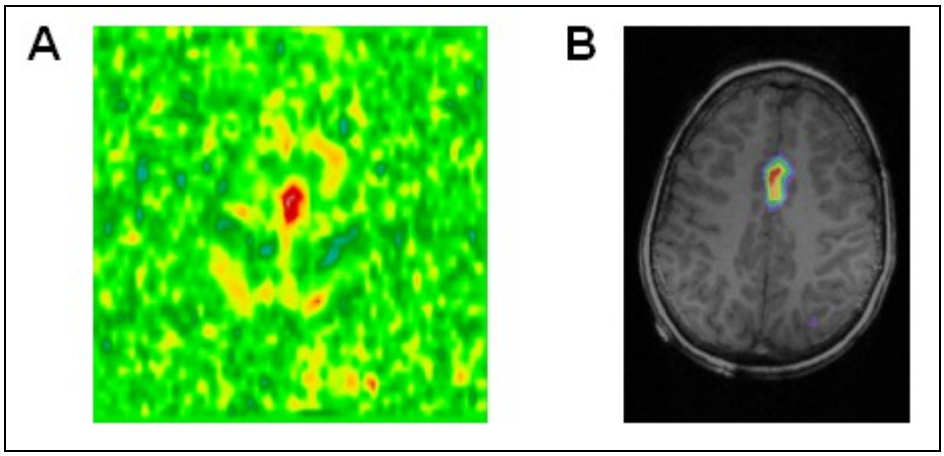

Figure 2.3-6 Example of (A) a raw t-map, and (B) the same t-map thresholded and overlaid on an anatomical MR scan.

2.3.6 fMRI Noise Sources and Pre-Processing

2.3.6.1 Motion

Although the subject's head is often secured with padding to minimize head movement during an fMRI scan, slight movements on the order of a few mm can occur and have an effect on voxel time courses as the placement of a given voxel within the brain slightly shifts. In general, an fMRI scan is pre-processed by use of a motion correction algorithm that realigns all the volumes recorded in a single run. Multiple runs from the same recording session can then also be realigned to one another if necessary. This realignment most often consists of adjusting for rigid body movement (i.e. translation and rotation of the head) and can only be effectively applied if there is only a few mm of movement. However, even with rigid body motion correction residual movement artifacts may exist.

For this reason, the six movement time courses calculated during the motion correction process (i.e. those describing translational and rotational movements in the x, y, and z directions) are often added as confounds to the design matrix used in the GLM. While these measures help minimize the effect of movement from one volume to another, correcting for movement that may occur between slices, i.e. movement that occurs during the acquisition of one of the volumes in a run, is considerably more complex and generally not carried out.

2.3.6.2 Drift

Drift consists of low frequency noise that occurs in essentially all fMRI scans. It is generally attributed to hardware instabilities (Lund et al 2006), particularly as it has been observed in phantom scans (Lund & Larsson 1999). Drift can be removed in pre-processing of the data by high-pass filtering, however if this is done one must ensure that the stimulation paradigm, i.e. the input function described in section 2.3.5, does not create signals lower than the filter's cut off frequency. Otherwise, it is also common to remove drift by including drift terms (e.g. cosines, polynomials, splines) that can model the drift seen in a given voxel time course so that it may be accounted for in the GLM (i.e. by adding the terms as confounds to the design matrix). For example, a 3rd order polynomial may be used to account for low frequency drift in a 6 minute long fMRI run (Worsley et al 2002).

2.3.6.3 Physiological Noise

Physiological noise can be generally attributed to two groups, respiratory and cardiac. Respiratory noise can be broken down into 3 components: 1) head movement resulting from respiration, which is actually reflected in the movement parameters, 2) susceptibility changes in the brain due to the movement of organs in the abdomen, and 3) respiration dependent changes in the oxygenation of the blood (Windischberger et al 2002). Although it has the largest effect on regions containing CSF pools (i.e. ventricles, outline of the

brain), its effects are seen throughout the brain (Windischberger et al 2002). Cardiac noise results from a person's heartbeat, which induces blood flow changes, and hence BOLD signal changes, in the blood vessels, while having little affect on other structures in the brain (Dagli et al 1999).

Unfortunately, while the temporal sampling frequency of a typical fMRI is around 0.5 Hz, the fundamental frequency of cardiac noise ranges between 0.6 and 1.5 Hz (Lund et al 2006). Likewise, it is not uncommon for a person to have a respiratory rate above 0.5 Hz while in the scanner, particularly if they are not used to the close quarters or scanner noises. Physiological noise is therefore heavily aliased during an fMRI scan and is often considered to be non-stationary (Lund et al 2006). In addition, while the effects of white noise diminish with higher magnetic field strengths, the effects of physiological noise actually intensify (Kruger & Glover 2001). All these characteristics make identifying and accounting for physiological noise a non-trivial task. For example, one method that attempts to remove the effects of physiological noise uses the signals recorded during the fMRI scan by a respiratory belt and pulse-oximeter (Glover et al 2000). These signals can then be used to detect the phases of the respiratory and cardiac oscillations at each fMRI time sample so that the noise can be modeled as a set of sines and cosines representing the aliased frequencies of the respiratory and cardiac noise. In another method the noise in a voxel time course is accounted for by breaking down the time course into components via independent component analysis (ICA) and then removing the component that best describes the signals seen in the CSF pools and blood vessels (Perlbarg et al 2007).

2.3.6.4 Spatial Smoothing

Spatial smoothing is a common pre-processing step that often consists of applying a 3-D Gaussian kernel of appropriate size (e.g. a 6 mm full width half

maximum) such that the voxels in an fMRI scan are spatially blurred by a slight amount. While this reduces the spatial resolution of the data, it helps increase the SNR. This is because while signals of interest, i.e. signals caused by changes in physiology, are expected to have a slight spatial correlation, white noise that arises in the recording process does not. By applying a spatial filter one is essentially averaging the signal within a given kernel. In fact, a better SNR can be obtained by recording volumes at half the resolution but twice the speed so that two volumes could be averaged to obtain one (Buxton 2002); this technique is usually not used but rather a Gaussian smoothing kernel of appropriate size (e.g. with 6 mm full width half maximum) is applied which will also condition the data for further statistical inference (i.e. the use of random field theory for the multiple comparisons problem).

2.3.6.5 Slice Timing Correction

In an fMRI scan a single frame (volume) is acquired over a time of about 2 s. As many slices are separately acquired within that volume, each slice represents data from a slightly shifted time point within that 2 s. In some cases this discrepancy between sample times is accounted for by the GLM applied by creating a specific design matrix for each slice by sampling each of the regressors at slice specific time points (Worsley et al 2002). Otherwise, slice-timing correction can be carried out as a pre-processing step by temporally interpolating all voxel time courses, often through linear or sinc interpolation, and then resampling each at a common time point (for example the time at which the first slice is acquired) (Henson et al 1999).

2.4 Localization in Epilepsy Through EEG-fMRI

While EEG remains the gold standard in the diagnosis, classification, and localization of epilepsy, it lacks spatial resolution, particularly for pre-surgical localization of the epileptogenic zone. On the other hand fMRI provides a much better spatial resolution, but by itself it is insensitive to epileptic activity and is rather most commonly used to localize sensory, motor, or cognitive functions. This is because the active and control states associated with these forms of function are easily controlled by the experimenter, allowing the input function used in the GLM process to be easily defined prior to the actual scan. Unfortunately, by definition, active states associated with epilepsy, i.e. epileptic events, occur randomly, resulting in an input function whose time course cannot be controlled. However, it is possible to define the active state in a patient as the time at which an epileptic event is detected in their EEG, while the control state can be defined as when their EEG is at baseline. The idea then behind EEG-fMRI is to record a patient's EEG while they are in the MR scanner so that an appropriate input function, based on events seen in the EEG, can be found and used to localize activity, which is correlated to the EEG events, with higher spatial resolution in the fMRI data than in the EEG. However, as one can imagine, recording an EEG within an MR scanner has its own obstacles. Whereas, the EEG signals themselves are very small, their sensitivity to the changing magnetic fields in the scanner is very large. Nevertheless, recording EEG activity of interest within this environment is possible (Ives et al 1993).

2.4.1 Issues with Recording EEG in an MR Scanner

One issue that arises with recording EEG in an MR scanner is the potential for the electrodes to burn the patient's scalp. As the electrodes are metallic, the changing magnetic fields of the scanner can cause the induction of fast currents that could potentially heat up the electrodes to the point that they would burn

the patient. This issue however can be avoided by using nonferrous electrodes, limiting resistors (Lemieux et al 1997), and by preventing the creation of current loops that involve the patient (Lazeyras et al 2001). Another issue that arises is the distortion of the EEG signal caused by slight movements of the head and wires in the static field of the scanner. This issue can easily be minimized through use of a very careful setup in which the head is immobilized by a vacuum cushion, the wires close to the head by bandages, and the wires leading to the amplifier by sandbags.

Another more serious issue that arises is the creation of artifacts in the EEG by the changing gradient fields of the scanner. These artifacts can be up to 50 times larger than the background EEG (Gotman et al 2006) and therefore an amplifier that has a high dynamic range, such that it does not saturate, must be used. This amplifier must then be connected to a computer outside the scanner room via a fiber optic cable to ensure that the scanner room's magnetic shielding is not broken. This gradient artifact can then be removed offline from the EEG signal assuming the EEG has been sampled at an adequate frequency (on the order of several kHz) (Benar et al 2003). The most common technique to do so consists of estimating the artifact, subtracting it from each frame, and then applying adaptive noise cancellation (Allen et al 2000). Another artifact commonly seen in EEGs recorded in an MR scanner is the ballistocardiogram artifact. This artifact is induced by small movements of the electrodes that result from the fast movement of blood in the arteries following each heartbeat. The ballistocardiogram artifact can be removed in a number of ways, some of which include averaging and subtraction, adaptive filtering, and through the use of ICA (Benar et al 2003).

While EEG recorded in an MR scanner suffers from many issues, good quality MR images can be obtained despite the presence of EEG equipment (Krakow et al 2000).

2.4.2 Data Acquisition and Analysis

The first experiments with EEG-fMRI used a technique of image acquisition referred to as "EEG-Triggered scanning". In this method, EEG artifacts induced by the scanner environment are avoided by recording a single fMRI volume after an event has been detected in the EEG (i.e. four or five seconds after a spike to catch the associated peak in the BOLD signal) (Gotman et al 2006). This volume would be deemed as representative of the active state and would then be compared to a control state volume to determine where the associated activation occurred. Drawbacks to this method include that only as many active volumes as detected events can be recorded, low-frequency drift cannot be accounted for, and an experienced electroencephalographer needs to be constantly monitoring the EEG during the scanning session.

As touched upon in section 2.4.1, the most common method of data acquisition in current use consists of continuously recording both EEG and fMRI data during a session and then removing artifacts in the EEG offline. After this has been done, an electroencephalographer marks the epileptic event times seen in the EEG (Figure 2.4-1A). These event times will then be used to create the input function, which will be convolved with an HRF to obtain the BOLD signal expected (Figure 2.4-1B) to result from the given EEG detected activity. Like most other fMRI studies, this expected signal is then used in the GLM process to create a t-map that defines which voxels are believed to be associated with the given activity (Figure 2.4-1D). As it has been found that differences in the HRF can exist from patient to patient, from one brain region to another, and even from one scanning session to another (Aguirre et al 1998), it is not uncommon to

carry out the GLM process multiple times with a series of different HRFs, each with a different delay in their peak time. The results of these processes can then be compiled to obtain a single t-map that describes the region of activation associated with the EEG detected activity.

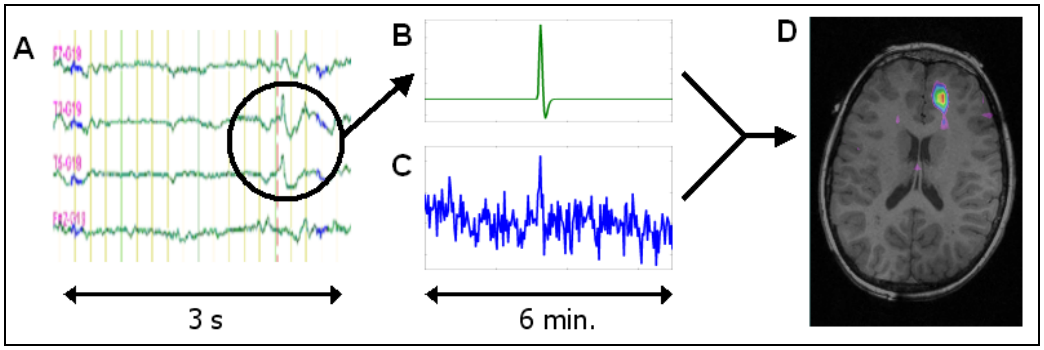

Figure 2.4-1 Example of EEG-fMRI data analysis process: (A) marking of EEG activity, (B) BOLD signal expected to result from marked EEG activity, (C) time course of a voxel involved with the detected activity, (D) thresholded t-map showing activated voxels.

2.4.3 Pros and Cons

In addition to maintaining a high temporal resolution thanks to the EEG, by incorporating fMRI into the recording process EEG-fMRI achieves high spatial resolution and allows BOLD signals related to epileptic activity to be recorded throughout the brain. Epileptic activity that is restricted to deep brain structures would not be detected however, due to the fact that the activity must first be detected in the EEG; this deep activity would be detected only if it was not restricted to deep brain structures and was actually correlated to superficial activity detected in the EEG. Recording EEG-fMRI is also somewhat cumbersome as it requires a very specific and delicate setup, particularly for acceptable recording of the EEG (see section 2.4.1), the disruption of which would lead unusable data. In addition, data analysis of EEG-fMRI data is dependent on adequate markings of epileptic events within the EEG, something that can vary from one electroencephalographer to another (Noachtar et al 1999).

2.5 Temporal Cluster Analysis

fMRI data consists of tens of thousands of voxel time courses. This means that, unlike EEG, it is impractical to detect activity of interest through visual analysis of the time courses. Temporal cluster analysis (TCA) is a data based fMRI analysis technique which was developed to map dynamic activity in the brain when the timing of said activity is not known. This technique works by exploiting the idea of temporal parcellation within the brain (Liu et al 1999), i.e. the probability of a voxel reaching its maximum value is equal and independent of time unless a stimulus is introduced. Put generally, TCA functions by collapsing the 4D data obtained in an fMRI scan into a 1D time course, or histogram, that consists of a count of the number of voxels, N_{max}, in the original data that achieve their maximum value at each time point.

2.5.1 Initial Studies Using Temporal Cluster Analysis

TCA was first developed and implemented in a study that hoped to determine when and where the satiation signal in human beings occurs (Liu et al 2000). After a person eats, the brain senses a biochemical signal and then sends a neurological signal indicating satiation. Unfortunately, the timing of this satiation signal is unknown, which in terms of fMRI data analysis amounts to not knowing the time course of the input function to be used in the GLM. TCA was therefore employed to detect when this satiation signal was occurring.

Subjects in the study ingested glucose at the 10 minute mark of a 48 minute fMRI scan. After the scan was completed, the data was processed offline using TCA to determine event times of interest. This consisted of finding the 1D histogram by looking at each time point in the scan and counting the number of voxels that reached their maximum value at that time point. Peaks in this histogram were then considered as events of interest. Figure 2.5-1 shows an example of the histogram found for one of the scans. As seen in this example, it

was found that TCA produced two peaks, one which occurred immediately after glucose ingestion and one that occurred 10.3 minutes after ingestion; this second peak was believed to result from the satiation signal. In contrast, no distinct peaks were seen for scans of control subjects who did not ingest glucose. Using these events times, appropriate input functions were created so that the activity associated with each peak could be mapped (activity associated with initial peak was found in the sensorimotor cortex, while activity associated with the second peak was found in the hypothalamus). It should be noted however that in this preliminary study only a single 10 mm thick mid-sagittal slice was imaged and processed.

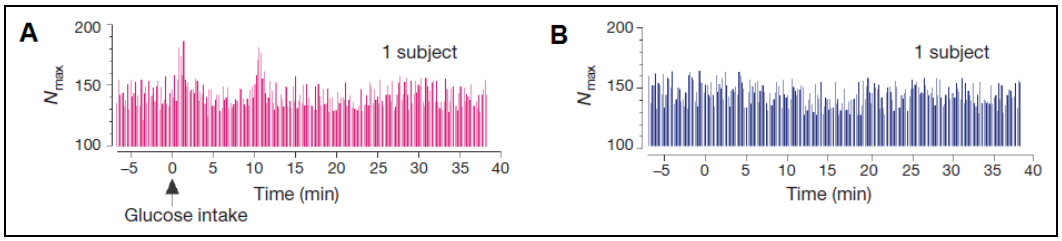

Figure 2.5-1 Example *1D histograms* found by TCA in (A) a subject who ingested glucose at time 0, and (B) a control subject. Adapted from (Liu et al 2000).

The application of TCA, although with slight variations, has been seen in a few studies following this initial investigation. In one, referred to as "weighted one-dimensional" TCA, rather than simply counting the number of voxels reaching their maximum value at each time point, the maximum values themselves were summed to create those found in *the histogram* (Yee & Gao 2002). In another study, TCA was carried out as an iterative process, each time on a smaller data set consisting of those voxels not counted in the largest peak of the histogram created in the previous iteration (Gao & Yee 2003). It has also been shown by one study (Zhao et al 2004) that TCA performs similarly to ICA in that they both generate similar activation maps for a set of event related fMRI experiments. However, TCA was found to be more computationally efficient and to have better repeatability.

2.5.2 Application of TCA to Epilepsy Localization

As the goal of TCA is to detect events of unknown timing, it is understandable why its application to the detection of epileptic activity in fMRI data, independently of the EEG, may be fruitful. This is particularly true due to the high temporal synchronicity associated with epileptic activity. If shown to be effective, TCA could potentially be used as a good alternative to EEG-fMRI for the localization of BOLD correlates to epileptic activity as it is less cumbersome, i.e. does not require the EEG, and could potentially detect activity restricted to deep brain structures (not visible on scalp EEG).

TCA was first used to localize interictal activity in fMRI data in 2004 by Morgan et al. (Morgan et al 2004). In this study, TCA was carried out on the fMRI data of two groups. The first group consisted of 6 TLE patients who had undergone resective surgery with successful seizure control, while the second group consisted of 3 patients whose seizure localization was deemed unclear in pre-surgical evaluations. The TCA technique carried out differed slightly from previous implementations in that only those voxels whose maximum value was between 2 and 10% above its first value, and above a background threshold, were counted to create the histogram. Histograms were then created for each run with peaks in the histogram defined as occurring at those time points when the histogram went above a value of 100 voxels. These peak times were then considered as defining periods of activation in an input function which was then passed to the GLM to obtain maps that localized the detected activity (see Figure 2.5-2). In the first group, higher t-values were consistently found in the hippocampus of the epileptogenic hemisphere when compared to those found in the hippocampus of the opposite hemisphere. In 3 of these cases statistically significant activation was seen in the appropriate hippocampus. Analyses involving the second group showed good agreement with the believed localization of the epileptogenic zone. The study concluded that the application

of TCA to epilepsy localization requires further testing, particularly on modifications to the technique that may provide better results.

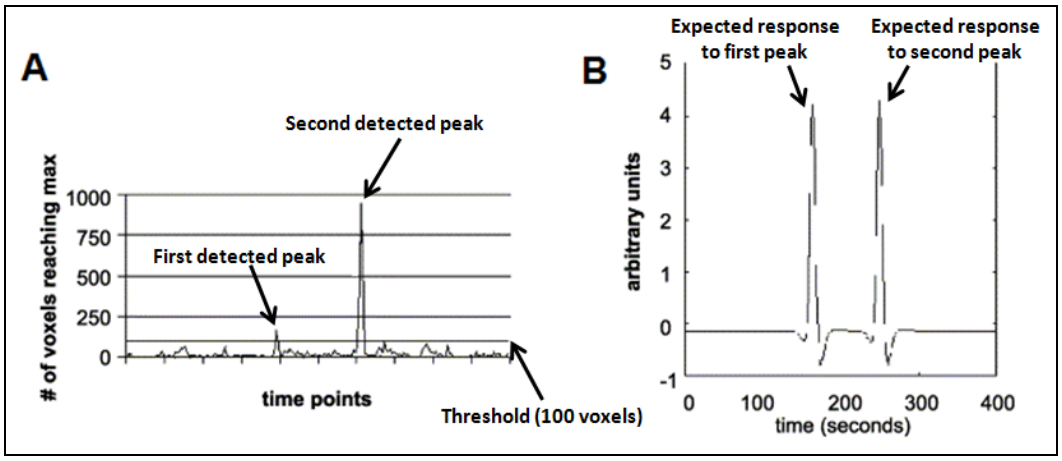

Figure 2.5-2 (A) Example histogram created and peaks detected by TCA and (B) expected BOLD response associated with the detected peaks. Adapted from (Morgan et al 2004).

2.5.3 Development of 2D-TCA for Epilepsy Localization

It has been found that BOLD signal changes due to epileptic discharges may be relatively small when compared to the influences of noise, motion, or other neuronal activity (Morgan et al 2008). Likewise, it was found that TCA is very sensitive to motion and physiological noise (Hamandi et al 2005). This is due to the fact that in TCA, as a single histogram is created, one is unable to differentiate between the different forms of activity that occur during a given scan. 2D-TCA is a technique that was developed to overcome this issue by creating multiple histograms, or reference time courses (RTCs), based on the timing of activity (Morgan et al 2008). Each of these RTCs may then be assumed to result from a different source.

2D-TCA consists of first detecting when events occur in a given voxel time course (i.e. whenever the value goes over a given threshold) to create an event time course (i.e. with a value of 1 where the events were detected and 0 elsewhere) for that voxel. This event time course is then added to one row of a 2D

histogram, each row of which consists of a separate 1D histogram. Initially, the time point at which the first event in an event time course occurred was used to define to which row the event time course was added (Morgan et al 2008); however, it was found that the peak time of the voxel may be a more appropriate characteristic to determine placement (Morgan & Gore 2009). For example, if voxel A and voxel B both reach their maximum values in frame 15 (i.e. volume 15 of the scan), their corresponding event time courses will be added to row 15, while if voxel C and voxel D reach their maximums in frame 77, their event time courses will be added to row 77. After this process has been completed, those rows, or RTCs, that show the largest amount of activation (i.e. have the largest event counts) are selected, grouped if found to be similar, and separately passed as input functions to the GLM so that corresponding t-maps may be obtained for each.

The initial study that implemented 2D-TCA (Morgan et al 2008) tested its application on computer simulated phantom data containing activity of amplitude between 1 and 4% above baseline, as well as on TLE patient data. Overall, when applied to the phantom data, the performance of 2D-TCA in producing accurate t-maps was better than that of TCA and comparable to the performance of ICA. This being said, 2D-TCA, which generally produces fewer components, was found to perform slightly worse than ICA when detecting multiple independent signal time courses. When applied to the small set of TLE patients, 2D-TCA produced t-maps showing mesial temporal activation, i.e. the expected area of activation. In addition, even though the patients were asked to remain still with their eyes closed, it was not uncommon that t-maps showing visual or motor activity as well as activity in areas associated with the default mode network (Gotman et al 2005) were found. In another study (Morgan & Gore 2009), the capability of 2D-TCA to detect and localize transient visual, auditory, and motor activity in control subjects was found to be similar to what

was achievable using event related processes (i.e. by use of a previously known input function), while in another study (Morgan et al 2010) 2D-TCA generally found the expected epileptogenic region in a homogenous group of 5 TLE patients.

2D-TCA has been shown to have the potential to detect and localize transient neuronal activity of unknown timing. This would be particularly useful in the localization of epileptic activity as it would do so with high resolution, could potentially detect activity restricted to deep brain structures, and would not be as cumbersome as EEG-fMRI. Unfortunately, the limits of its capabilities to detect various forms of activity in terms of activation size, amplitude, and frequency have not been evaluated. This is an important step in determining whether or not the general application of 2D-TCA to detect epileptic activity is worthwhile.

3 Methods

3.1 Data Acquisition

All data was obtained from a database of individuals (patients with epilepsy and control subjects) who underwent EEG-fMRI acquisitions. Images were acquired with a 3T Siemens Trio Scanner. Anatomical MR images were acquired with the following parameters: TR = 23 ms, TE = 7.4 ms, flip angle of 30^{o}, 1 mm isotropic voxel size, 256x256 matrix, 176 sagittal slices. Functional MR images were acquired over 6-min scanning runs with the following parameters: TR = 1.75s, TE = 30 ms, flip angle of 90^{o}, 5 mm isotropic voxel size, 64x64 matrix, 25 transverse slices.

Runs were acquired over a 2 hour scanning session unless the subjects were prematurely taken out of the scanner due to discomfort or, in the case of patients, if they had a seizure during the scan that could have potentially caused them injury. A single scanning session then consisted first of acquisition of the anatomical MRI, followed by acquisition of approximately 9 fMRI runs (sometimes more, sometimes less depending on time constraints). EEG was continuously recorded during this time with 25 MR-compatible Ag/AgCl electrodes placed on the scalp according to the 10-20 system (19 standard locations) referenced to FCz with extra electrodes at F9, F10, T9, T10, P9, and P10. In addition, two electrodes were placed on the upper back to record the electrocardiogram. All electrodes were adjusted such that their impedance was below 5 kΩ, a value that was monitored throughout the scanning session so that electrodes could be readjusted between runs if needed. The EEG was low pass filtered at 1 kHz and sampled at 5 kHz using a BrainAmp amplifier (Brain Products, Gilching, Germany). After scanning had completed, gradient artifacts were removed from the EEG using an averaged subtraction method (Allen et al 2000) implemented by BrainVision Analyzer software (Brain Products, Gilching,

Germany). The ballistocardiogram artifact was then removed using an ICA method (Benar et al 2003). An electroencephalographer would then visually inspect the EEG and mark the timing of any recorded epileptic activity.

3.2 Simulated Data

A large set of fMRI scans containing simulated epileptic activity was created by adding BOLD signals, simulated using values based on current knowledge of BOLD responses to epileptic activity, to time courses of voxels in specific regions of interest (ROIs) in 6 control subject runs (see Figure 3.2-1). Simulated responses were created by convolving a Glover HRF, of appropriate amplitude, with a simulated epileptic activity time course (i.e. a function that had a value of 1 when activity occurred and 0 elsewhere). Activity associated with spikes was simulated in 3 ROIs (left temporal lobe, right frontal lobe, and right hippocampus) using all combinations of the following characteristics (all values used for a given characteristic were chosen based on values commonly seen in the EEG-fMRI results of epileptic patients): 1, 5, or 10 randomly timed spikes per run; HRF amplitudes of 0.5-2% above baseline, in 0.25% increments; and ROI sizes of 12, 27, 36, 64, 80, and 125 voxels. In addition to spikes, a single 5 s long event was simulated in a parietal lobe ROI with the above HRF amplitudes in ROI sizes of 64, 125, and 216 voxels.

Each run was simulated such that it contained all 4 forms of epileptic activity (i.e. 1 spike, 5 spikes, 10 spikes, and one 5 s event), one in each ROI. For example, a given run may contain 1 spike in the temporal lobe, 5 spikes in the frontal lobe, 10 spikes in the right hippocampus, and a 5 s event in the parietal lobe. Simulations were repeated 6 times such that each of the three ROIs containing spikes could simulate each number of spikes (1, 5, or 10) twice; accordingly, simulations of the 5 s event were repeated 6 times. This allowed for the final

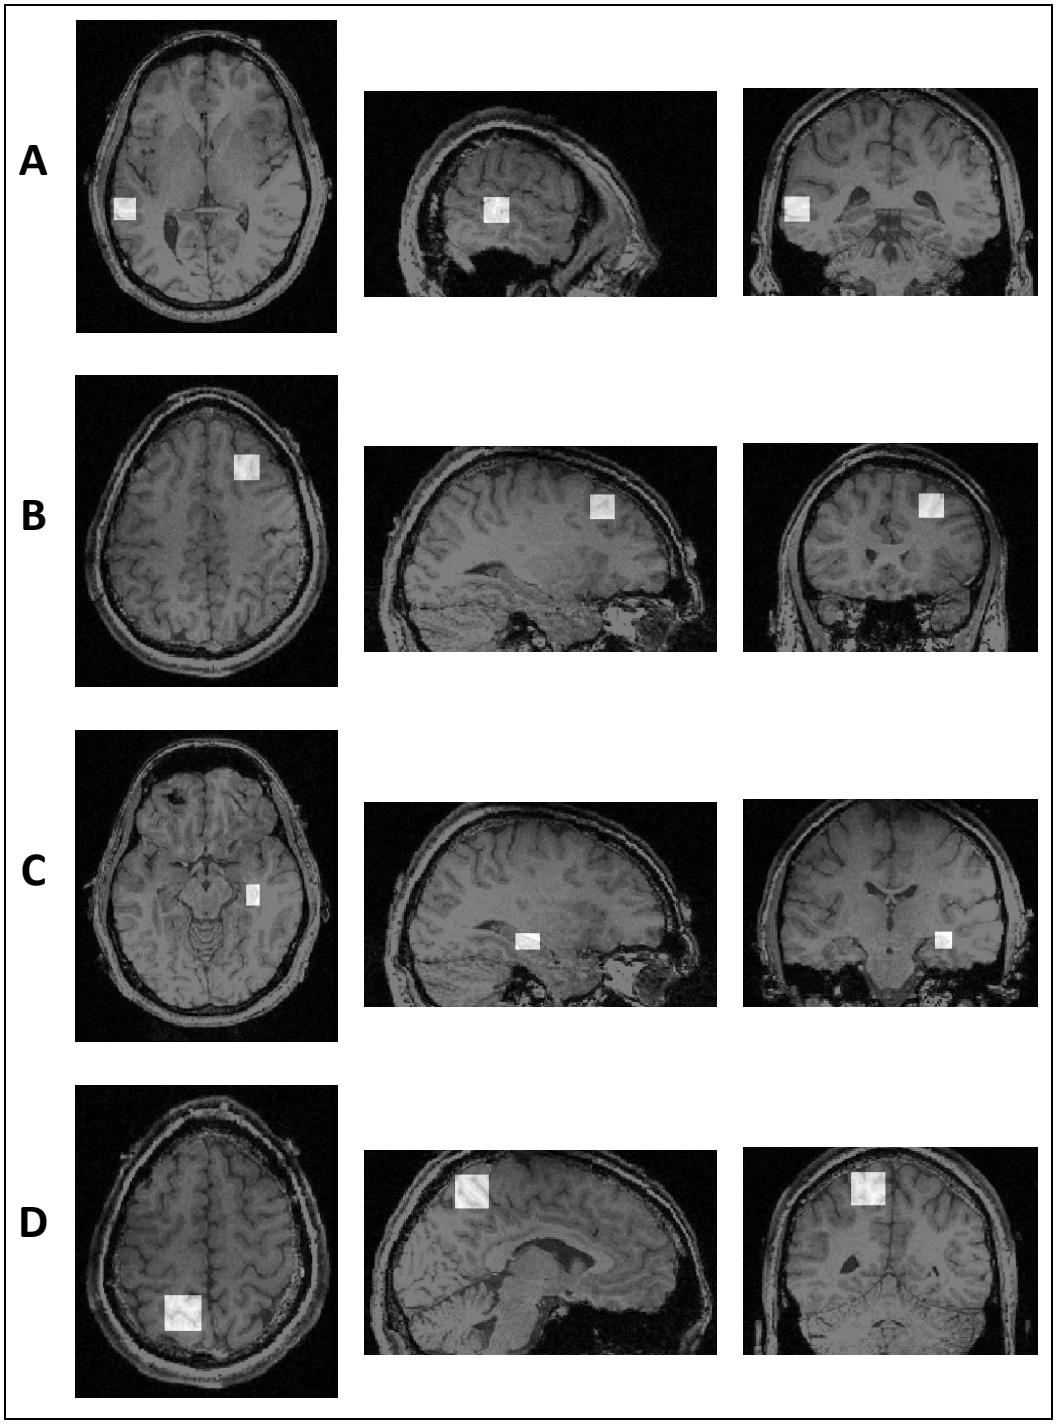

Figure 3.2-1 Examples of 4 ROIs to which simulated data was added: (A) left temporal lobe (3x3x3 voxels), (B) right frontal lobe (3x3x3 voxels), (C) right hippocampus (2x2x3 voxels), (D) left parietal lobe (4x4x4 voxels). In addition to the cluster sizes shown, larger versions of each ROI were simulated by increasing each dimension by 1 voxel (i.e. 3x3x3 would become 4x4x4 and then 5x5x5).

number of uniquely simulated runs to be 756 (6 control subject scans x 3 ROI sizes x 7 HRF amplitudes x 6 repetitions), for a total of 3024 unique simulated BOLD responses to epileptic activity (4 forms of activity/run x 756 runs). After being created, all simulated runs were pre-processed (see section 3.3.1). Then, using these runs, a modified version of the 2D-TCA algorithm as developed by Morgan and Gore (Morgan & Gore 2009) was created and its limits, in terms of responses that it is able to detect, were investigated.

3.3 Steps/Development of the 2D-TCA Algorithm

3.3.1 Pre-Processing

Pre-processing consists of 4 main steps. First, as the nature of 2D-TCA is to aggregate events seen in voxels that peak at the same time, it is important to carry out slice timing correction to ensure that voxel time courses from different slices are re-sampled at the same time points. This is followed by within run motion correction and spatial smoothing by a 6 mm full width half maximum Gaussian kernel. Finally, using the associated anatomical scan, an appropriate brain mask is created for the functional data so that only those voxels recorded within the brain are considered in the following steps; this differs from the technique implemented in (Morgan & Gore 2009) in that they simply apply a background threshold to differentiate between those voxels that should be considered and those that should be ignored.

3.3.2 Temporal Filtering

While in (Morgan & Gore 2009) high frequency white noise was removed from voxel time courses by applying a 3-point temporal averaging filter, it was felt that applying a rectangular window to a voxel's frequency spectrum, such that those frequencies that a BOLD response is not expected to contain would be removed, would be more appropriate as its effects are more easily understood.

In addition, the effects of low frequency drift could be removed by this means. As 2D-TCA is only interested in detecting activity associated with the HRF, it would be appropriate to define the window's borders according to the frequency spectrum of the HRF. Figure 3.3-1(A) shows the frequency spectrums of the Glover HRF, as well as the frequency spectrum of the Glover HRF after being convolved with an 8 s step function (i.e. 8 s of simulated activity). The low cut-off frequency of the window can then be defined by the low -3 dB mark on the 8 s event spectrum (i.e. at 0.006 Hz), while the high cut-off frequency can be defined by the high -3 dB mark on the HRF spectrum (i.e. at 0.085 Hz), as shown in Figure 3.3-1(A). Figure 3.3-1(B) shows an example of a raw voxel time course that contains 5 spikes (indicated by the red impulses), simulated using an HRF amplitude of 1% above baseline, while Figure 3.3-1(C) shows the same voxel time course after frequency components not within the designated window have been removed. It should be noted that although this window is applied, as shown in Figure 3.3-1(C) the mean value of the time course (i.e. the 0 Hz frequency component) is not removed; this is so that the next step, which appropriately normalizes the data, can be effectively carried out. It should also be mentioned that in applying this frequency spectrum window, BOLD responses to epileptic events with duration longer than 8 s will be removed from voxel time courses and therefore are expected to be undetectable with this 2D-TCA technique.

3.3.3 Baseline Definition and Normalization

BOLD signal data is recorded in arbitrary units. Therefore, to compare activity occurring in different voxels, their time courses must first be normalized to the same scale. Ideally, one would normalize a voxel by its baseline value. However, this would require one to know the timing of activity and inactivity, the very characteristic which 2D-TCA aims to discover. As such, the baseline value of a

given voxel time course must be estimated by other means. In (Morgan & Gore 2009) a voxel time course is normalized by the average of its first 5 values. This

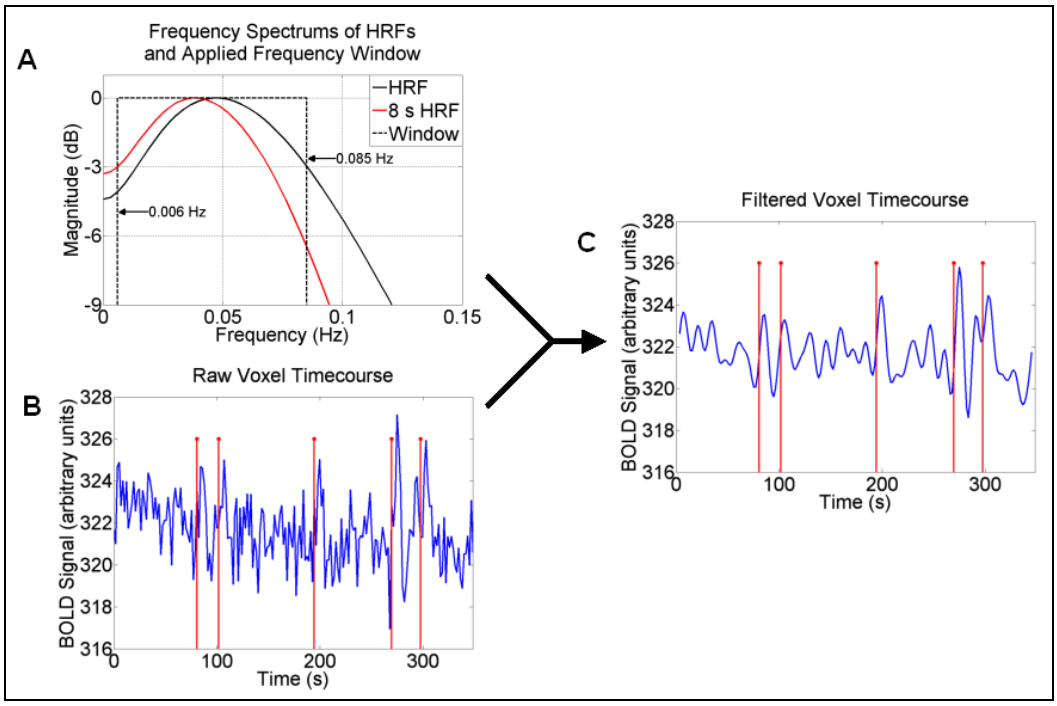

Figure 3.3-1 (A) Plots of HRF Frequency spectrums with indicated band pass filter cut-off frequencies. Also, an example of a voxel time course containing 5 simulated spikes (B) before the band pass filter is applied, and (C) after.

technique is not very robust as: 1) the average is calculated from only 5 values and therefore is not a particularly strong estimation of baseline in the statistical sense, and 2) it assumes that a patient is at baseline during the first 5 frames of the fMRI scan, a characteristic that cannot be controlled due to the intrinsic randomness of epileptic activity. Therefore, other methods were considered.

The first method consisted of defining a voxel's baseline value by the mean of its time course. Unfortunately, due to the non-symmetric nature of the HRF, using the mean of a voxel time course would theoretically provide a positively biased baseline value (see Figure 3.3-2(A)). As such, a second technique was employed. This consisted of applying k-means clustering analysis, a technique that aims to partition *n* observations into *k* clusters in which each observation belongs to the

cluster with the closest mean. To find an appropriate baseline, the n = 200 values (i.e. the number of time samples) of a voxel time course were grouped into k = 2 clusters, one containing the higher values seen in the voxel time course, assumed to represent values associated with activity in the voxel, and one containing the lower values, assumed to represent values associated with inactivity. The mean of the cluster that contained the lower values would then be used to estimate the baseline value for that voxel. While this technique may theoretically impose a slight negative bias on the baseline value (see Figure 3.3-2(A)), it is more effective than simply using the mean. Figure 3.3-2(B) shows the baseline value (black line) estimated using the k-means technique for the same voxel time course as shown in Figure 3.3-1(C) (for comparison, the green line indicates the mean value); Figure 3.3-2(C) shows this same time course after it has been normalized by the k-means baseline value.

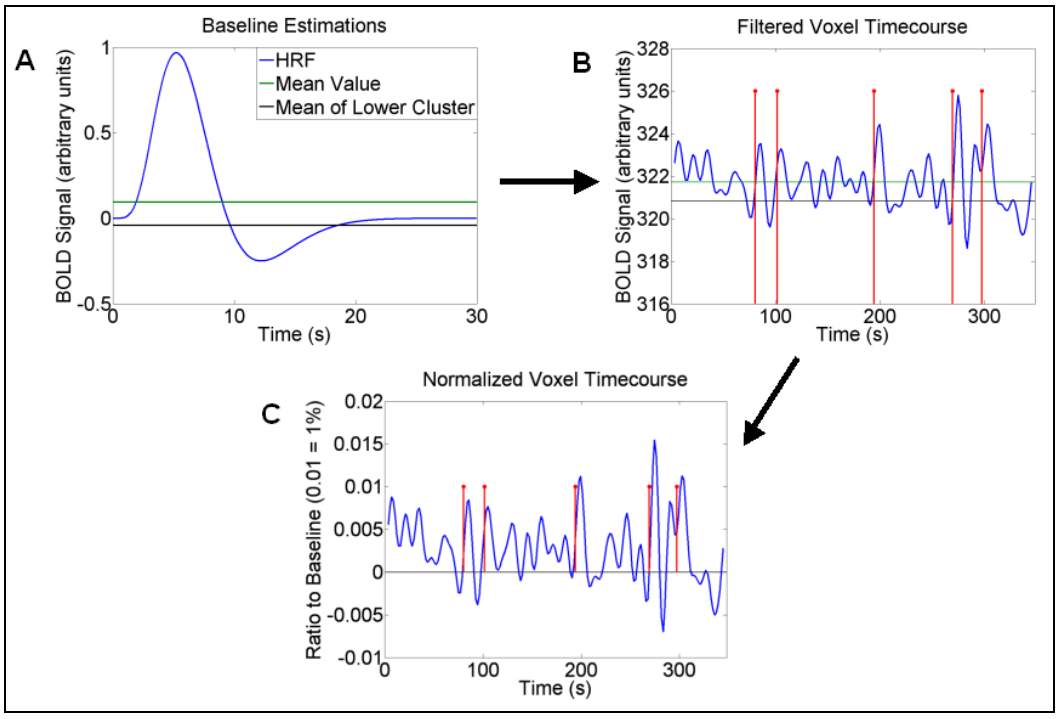

Figure 3.3-2 (A) Estimates of baseline for the Glover HRF using the mean and through use of the k-means clustering process. (B) Application of the k-means process to define baseline value in an example time course (mean value is also shown in green). (C) Normalization of the time course by the k-means calculated baseline.

3.3.4 Detection of Candidate Voxels

Having brought all voxel time courses onto the same scale, those voxels that may show activity of interest must be differentiated from those that do not. In (Morgan & Gore 2009) this is carried out by defining a predetermined threshold range, between 0.5 and 8% above baseline, within which BOLD signal changes associated with the activity of interest are expected to exist. While the lower threshold of 0.5% was implemented to ignore those voxels which can be considered to be inactive during the run, the upper threshold of 8% allows 2D-TCA to ignore voxels whose activity may be much larger than what would ever be expected to occur as a result of neural activity (i.e. activity larger than 8% may be attributed to some form of noise). If a voxel's maximum value exists within this range, the voxel is then considered as a "candidate".

For this study, it was decided that a similar technique would be employed, but that a number of range boundaries would be tested to find the most appropriate values. Lower boundary values tested were chosen to border the HRF amplitudes used to create the simulated data (i.e. the values of BOLD responses expected to be associated with epileptic activity); values tested were 0, 0.5, 1, 1.5, and 2% above baseline. On the other hand, upper boundary values tested were chosen to border the 8% value implemented in (Morgan et al 2008); values tested were 3-11% above baseline with increments of 1%. All combinations of range boundaries were tested on all simulated runs to determine their ability to differentiate between those voxels that contained simulated activity and those that did not. The best range was determined as that which provided the best average specificity across runs while maintaining an average sensitivity of at least 90% to the voxels that contained simulated epileptic activity. This range was found to be between 1 and 6% above baseline and provided a true positive rate (TPR) of 0.90 and a false positive rate (FPR) of 0.59. Figure 3.3-3 shows the

application of this threshold to the same example voxel time course as in Figure 3.3-2(C).

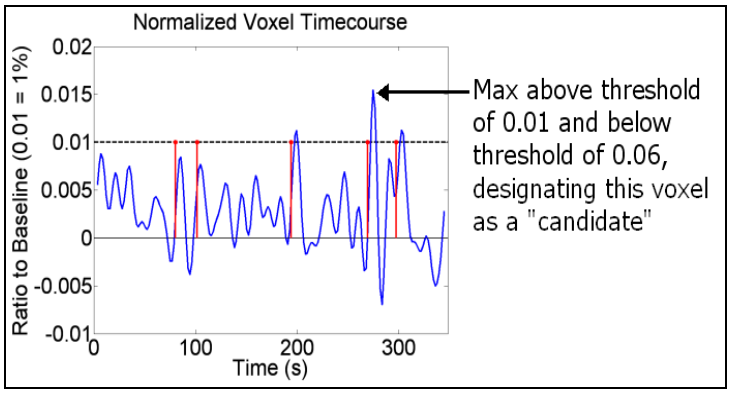

Figure 3.3-3 Example of a voxel being defined as a candidate based on its maximum value.

3.3.5 Global Time Course Removal

After designating which voxels are candidates, the 2D-TCA technique employed in (Morgan & Gore 2009) creates a “global voxel time course”, calculated as the mean of the time courses of those voxels that are not candidates. This global voxel time course is then subtracted from all candidate voxel time courses. Although, the exact effect of this may not be clear, it is expected that doing so removes any global forms of activity or noise that may mask the BOLD signals specific to each candidate voxel. Although no in-depth investigation was carried out, this study found that removing this global time course did in fact allow for somewhat better performance in the next step, event detection within the voxel (section 3.3.6). For this reason, calculation and removal of the global time course was retained as a step in the 2D-TCA algorithm developed for this study. Figure 3.3-4(A) shows an example of a global voxel time course, while Figure 3.3-4(B) shows the time course from Figure 3.3-3 after this global voxel time course has been removed. In this example it can be seen that removal of the global time course has a beneficial effect in that through this process activity seen in the first 50 seconds of the voxel’s time courses is reduced, while activity associated with the first and second spikes is increased.

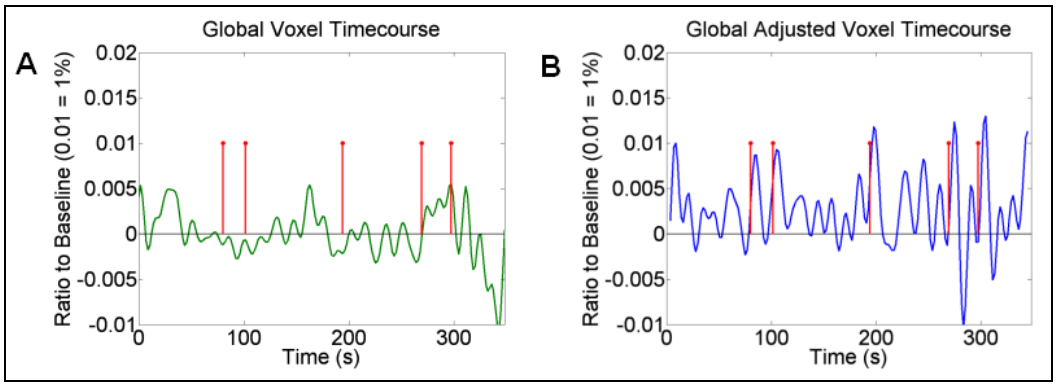

Figure 3.3-4 Examples of (A) a global voxel time course, and (B) a voxel time course after the global voxel time course has been removed.

3.3.6 Event Detection within a Voxel

To detect when events occur in a given voxel time course the technique in (Morgan & Gore 2009) applies a threshold of 1.5 standard deviations, i.e. any time the voxel's value is larger than 1.5 standard deviations above the mean of its time course, an event is counted. It is expected that values in the voxel time course which peak by a significant amount above the voxel's standard deviation indicate points of transient neural activity, similar to what is expected to result from epileptic activity. This technique was retained in the 2D-TCA algorithm developed for this study. However, as the definition of baseline in this study differs from that in (Morgan & Gore 2009) (see section 3.3.3), a different, more appropriate, threshold value needed to be found. A variety of threshold values above baseline, all multiples of the voxel's standard deviation, were tested on every simulated voxel time course to determine which provided the most accurate event detection. A total of 11 threshold values were tested ranging from 0 to 2 standard deviations above baseline in increments of 0.2. The best threshold value was determined as that which provided a maximum average specificity to events in all simulated time courses while maintaining an average sensitivity of at least 90% to those same events. This value was found to be 1.2 standard deviations above baseline and provided a TPR of 0.91 and an FPR of 0.29. Figure 3.3-5(A) shows the application of this threshold to detect events in

the time course from Figure 3.3-4(B), while Figure 3.3-5(B) shows the resulting event time course for the voxel.

In addition, although not implemented in (Morgan & Gore 2009), it was decided that a spatial constraint would be imposed in the event detection step due to the high level of synchronicity and spatial clustering associated with epileptic activity. This constraint consisted of counting an event in a voxel time course only when events were also detected at the exact same time point in 4 of the voxel's 6 closest neighbours. By incorporating this condition it is hoped that those events that may occur due to white noise, or other activity that does not have a strong spatial correlation, would be ignored. Although no detailed investigation was done, this criterion was chosen on the basis that it did not seem overly conservative or liberal. In actuality, it is very likely that better criteria may exist, however, this would require deeper investigation.

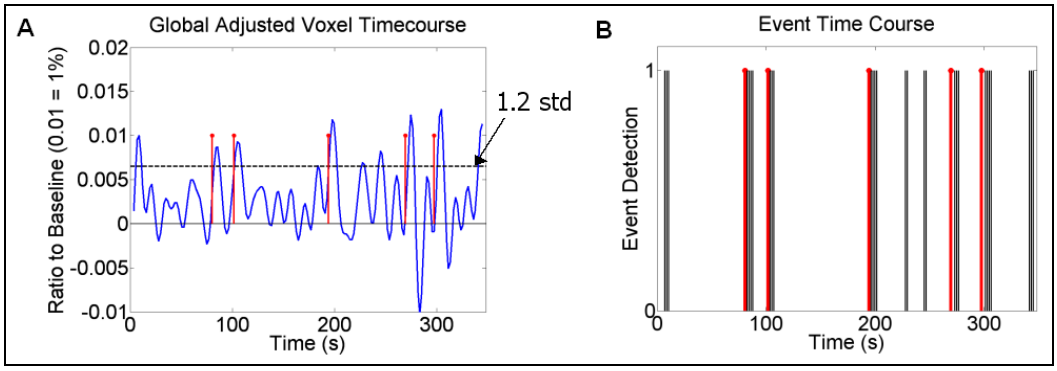

Figure 3.3-5 (A) Event threshold being applied to an example voxel time course, and (B) the resulting event time course.

3.3.7 Creation of 2D Histogram

Creation of the 2D histogram is a crucial step in 2D-TCA as this is where the actual temporal clustering is carried out. The technique described in (Morgan & Gore 2009) aggregates the event time courses of different candidate voxels into a 2D histogram according to the time frame in which the voxels reach their maximum values, i.e. the frame in which a voxel reaches its maximum value

defines to which row in the 2D histogram its corresponding event time course will be added. This similarity mechanism makes sense as it is expected that the voxels associated with the same activity will generally peak at the same time. This event time course clustering technique was therefore retained in the 2D-TCA algorithm developed in this study. Figure 3.3-6 shows this clustering process being carried out on three example voxel time courses (voxel 1 being the voxel whose time course is shown in Figure 3.3-5(A)). As voxel 1 and voxel 2 both peak at frame 174 (Figure 3.3-6(A)), their event time courses (Figure 3.3-6(B)) will both be added to row 174 of the 2D histogram (Figure 3.3-6(C)). On the other hand, as voxel 3 peaks at frame 93, its event time course will be added to row 93 of the 2D histogram. After the 2D histogram has been created, those rows to which no event time course was added are discarded.

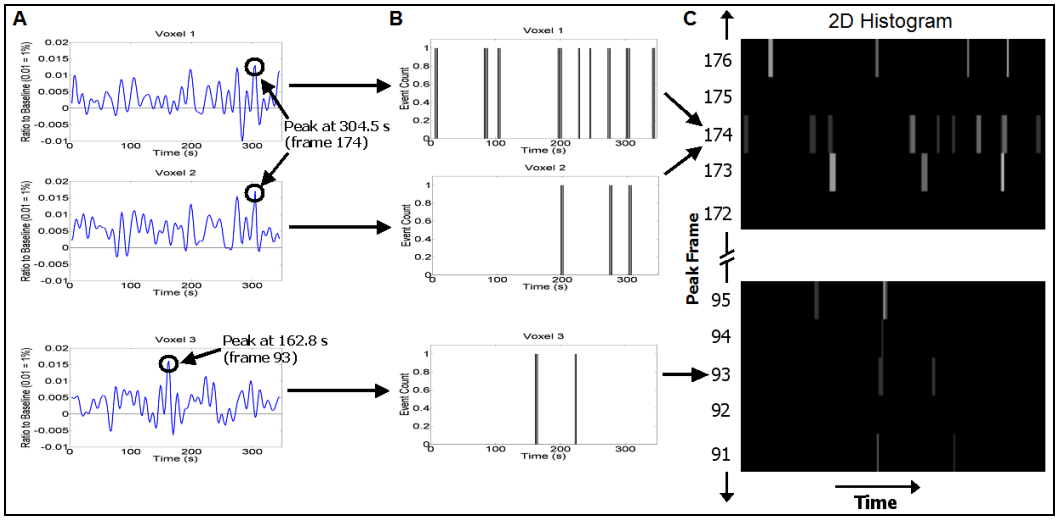

Figure 3.3-6 Process of 2D histogram creation showing (A) detection of voxel peaks, and addition of (B) event time courses to appropriate rows of (C) the 2D histogram (a lighter shade of gray indicates a higher event count at that time point in that row; black indicates an event count of 0).

3.3.8 Grouping Similar Rows in the 2D Histogram

The temporal clustering technique applied in the previous step aims to create many time courses, each of which describes, hopefully uniquely, underlying activity common to many voxels. However, it is probable that after the 2D

histogram is created, more than one row will describe the same activity. This arises from the fact that voxels associated with the same activity may in fact peak at different times due to slight variations in their time courses. Events in the time courses of these voxels will still be detected at similar times. As such, those rows of the histogram that describe similarly timed activity should be grouped.

In the 2D-TCA algorithm developed in this study, grouping was carried out in two steps. The first consisted of comparing the time courses of all rows to one another and then summing those whose correlation coefficient to each other was above a certain threshold. A range of correlation coefficient thresholds, from 0 to 1 in increments of 0.1, was tested on all simulated runs. To determine the performance of each threshold value, t-maps, thresholded at $t > 3.1$ ($P < 0.001$), were created from each of the grouped histograms that resulted from applying the given threshold value. Those t-maps whose regions of activation best described the four simulated ROIs in a given run, determined by their TPR (the associated FPR of voxels within the brain was generally negligible), were selected and the overall TPR associated with the given threshold value was calculated as the number of true positives in those selected t-maps divided by the number of voxels known to be active in that run. The best threshold was then chosen as that which performed the most grouping of rows (i.e. the lower the threshold the more the grouping) while maintaining a reasonable TPR. Figure 3.3-7 shows both the average TPR and number of resulting rows for all tested threshold values. As the TPR started to severely decline for a threshold of 0.7 and below, a threshold of 0.8 was found to be appropriate (i.e. if the correlation coefficient between two rows is 0.8 or higher they will be summed). This provided an average TPR across all simulated runs of 0.81 and an FPR, which for a given run was calculated as the average of the FPRs of the four selected t-maps, of 0.012; by applying this threshold an average of 140 rows would remain

after grouping had been carried out in this first step (for comparison, if no grouping is done an average of 183 rows would remain).

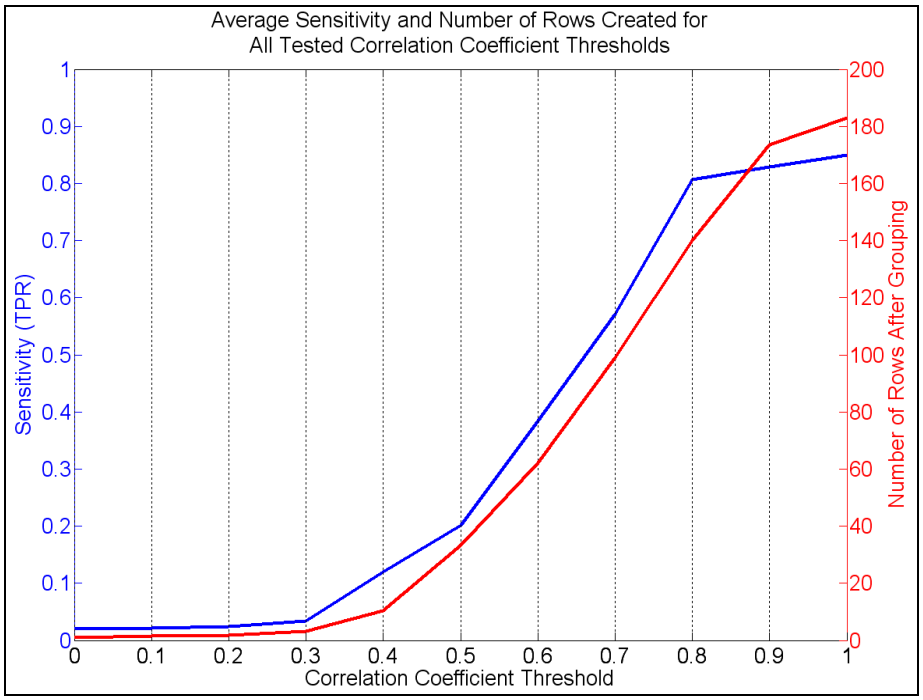

Figure 3.3-7 Average TPR and number of rows created for all tested correlation coefficient threshold values.

Unfortunately, simply grouping according to correlation is not sufficient as those rows that describe similar activity, but whose peaks occur at distant time points, will not be grouped due to the nature of the correlation coefficient calculation. The second step applied in the grouping technique aims to account for this situation. It consists of applying a threshold to each row, equal to the mean of the non-zero values of that row, and then grouping those rows whose activity above threshold overlaps for at least a certain percentage of their combined time of activity. The step would be performed after the previous grouping by correlation. A range of time overlap values, from 0 to 100% in increments of 10%, were tested on all simulated runs with performance being determined by the same technique applied to determine performance of the first grouping step. Figure 3.3-8 shows both the average TPR and number of resulting rows for all

tested threshold values. As the TPR started to severely decline for a threshold of 20% time overlap and below, a threshold of 30% was found to be appropriate (i.e. if the time courses of two rows, thresholded by their mean values, overlaps for more than 30% of their combined time of activity, they are grouped). This provided an average TPR across all simulated runs of 0.70 and an FPR of 0.026; by applying this threshold an average of 59 rows would remain after grouping had been carried out by this second step. The two above described grouping steps were not parts of the original 2D-TCA algorithm.

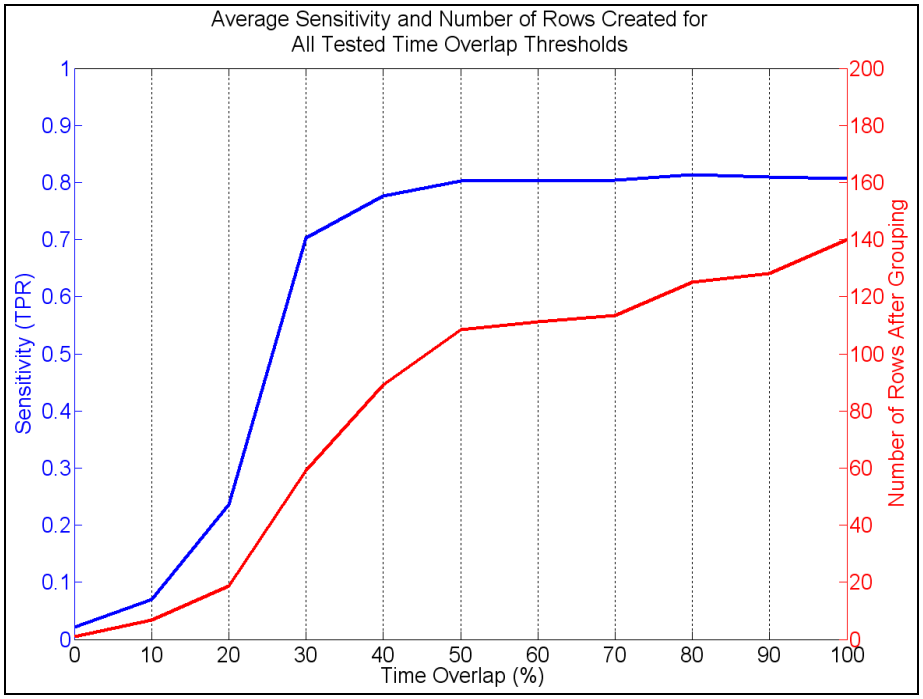

Figure 3.3-8 Average TPR and number of rows created for all tested time overlap threshold values.

Figure 3.3-9 shows examples of three rows whose time courses were found to be similar, as well as the final component resulting from the sum of these rows. While the time courses of row 130 and 131 are deemed similar through use of the correlation coefficient, the time course of 121 is not as its peak value occurs at a distant frame number. However, by applying the second grouping step the time course of row 121 is also considered as similar to those of row 130 and 131.

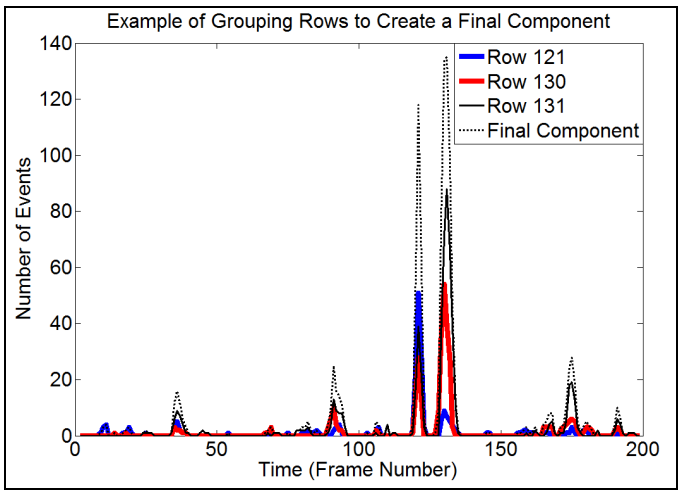

Figure 3.3-9 Examples of similar rows and the final component arising from their sum. While row 130 and 131 are deemed similar through correlation, row 121 is deemed similar by the second step in the grouping process.

3.3.9 Removal of Insignificant Components

After similar rows have been grouped, the resulting time courses can be considered as separate components because each will be passed to the GLM as a separate input function and will therefore have a corresponding t-map of activation. However, before this is done, to lower the number of t-maps that will be created only those components with significant activation will be passed to the GLM. In (Morgan & Gore 2009) the mean and standard deviation of component maximum values are calculated. Those components whose maximum value is larger than 1 standard deviation above the calculated mean are then considered as significant and are passed to the GLM. In applying this threshold, this study found that components describing simulated activity were often inappropriately discarded. As such, the 2D-TCA algorithm developed in this study takes a less conservative approach by applying a lower threshold, applied to the actual peak number of events in a component, such that those components that can be considered as insignificant are removed. A range of threshold values, from 0 to 65 in increments of 5, were tested on all simulated runs with performance being determined by the same technique applied to determine performance of the two grouping steps (see section 3.3.8). As this

measure of the TPR started to severely decline at a threshold of 20 events or higher, a threshold of 15 events was found to be appropriate (i.e. a component was discarded if its maximum event count was below 15). This provided an average TPR across all simulated runs of 0.65 and an FPR of 0.029. A threshold of 15 was chosen over lower values, which provided lower FPRs, as it produced fewer components (on average a total of 15). Figure 3.3-10 shows examples of a component that is deemed insignificant by this threshold and another that is retained.

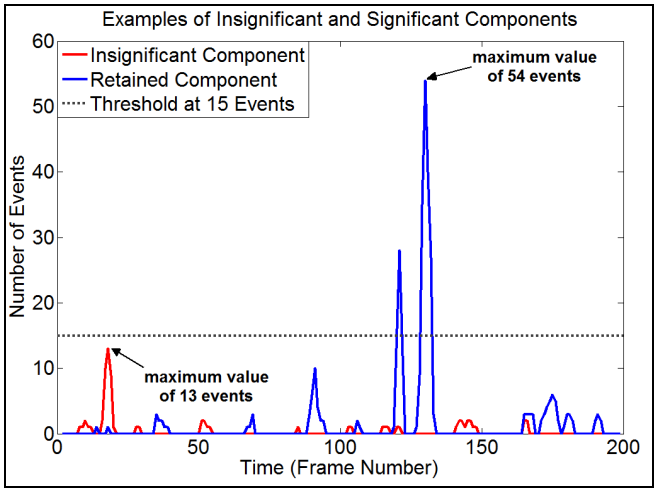

Figure 3.3-10 Examples of a component that is deemed insignificant and a component that is retained.

3.4 Creation of t-maps from Components

Components created by 2D-TCA are then passed, along with the motion parameters calculated during motion correction, to the GLM (Worsley et al 2002) as input functions to obtain their corresponding t-maps. These maps were then thresholded at $t > 3.1$ (corresponding to uncorrected $P < 0.001$) to determine regions of activation. It should be noted that as component time courses describe changes in the BOLD signal itself rather than neuronal activity, the peak value of the HRF used in the GLM does not include a delay. In fact, the HRF that was used in this study was simply a version of the Glover HRF, positioned such that the peak occurred at time 0.

3.5 Selection and Analysis of Patient Data

Patient data used to test the capabilities of 2D-TCA were selected based on the technique's performance in detecting the simulated activity. This consisted of determining for what characteristics (i.e. event frequency, HRF amplitude, and activation cluster size) the 2D-TCA algorithm was able to effectively detect simulated activity. Patient runs containing activity with similar characteristics, determined from EEG-fMRI results, would then be selected from a large database (only those patients who showed clear activation in their EEG-fMRI results were considered). This resulted in 2D-TCA being tested on a total of 60 runs obtained across 20 separate patients. In 40 of these runs interictal spikes were detected by the EEG, while in the other 20, prolonged forms of interictal activity were detected (nothing longer than 8 s was selected).

The effectiveness of 2D-TCA to detect epileptic activity within a given patient run was determined by qualitatively comparing its results to results obtained by EEG-fMRI for the same run. Cases in which 2D-TCA created at least one component (i.e. a t-map) that accurately described what was seen by EEG-fMRI were further investigated by a neurologist. The neurologist was asked to consider the results of 2D-TCA as a substitute for EEG-fMRI. This was done by blinding the neurologist to the EEG-fMRI results, but allowing them to have full access to all other patient data (i.e. routine EEG, anatomical MRI, clinical data, etc.). Given this knowledge, the neurologist was then asked to rank each of the components created by 2D-TCA based on how likely it is that they describe epileptic activity within the given patient. This was done on a scale of 1 to 5, 1 indicating a component that is definitely not associated with epileptic activity and can therefore be ignored, 3 indicating a component whose source is unclear (i.e. could be epileptic activity or not), and 5 indicating a component that is most likely arising from epileptic activity (scores of 2 and 4 were simply intermediary values used to give the neurologist some flexibility).

4 Results

4.1 Performance in Detecting Simulated Activity

Ideally, the only outputs created by 2D-TCA when it is applied to a simulated run would be t-maps that describe those ROIs containing simulated activity. However, it should be kept in mind that the final output of 2D-TCA will most likely also include other t-maps describing other forms of transient activity detected in a run. It is therefore important to determine whether or not within this larger set of generated t-maps there is one that corresponds to each of the four ROIs simulated in a run (i.e. a total of four t-maps, each describing a different form of simulated activity, i.e. 1 spike/run, 5 spikes/run, 10 spikes/run, or a 5 s event). Figure 4.1-1 shows the average TPR (calculated across all runs containing the specified form of simulated activity) for the t-map produced by 2D-TCA whose region of activation best described the given form of activity simulated in the run, i.e. that t-map which had the highest TPR for the associated ROI (the corresponding FPR values of voxels within the brain were all very small, on the order of 0.01). For each form of activity, this value is shown across all simulated HRF amplitudes and ROI sizes. While in all cases an increase in the HRF amplitude used to simulate the activity leads to an increase in the TPR, an increase in the ROI size has nearly no effect. In fact, no major trend associated with ROI size was expected to be seen, except that 2D-TCA may have the most trouble in detecting a 12-voxel ROI of activity as only those components with a maximum count of 15 voxels or more are retained to create corresponding t-maps. In fact, looking at Figure 4.1-1, one could argue that, when detecting spikes, 2D-TCA performs the worst when the simulated ROI is only 12 voxels large. We attribute the lack of a more drastic difference to the counting of non-simulated voxels that happen to peak at the same time point, which increases the maximum of the given component to 15 or more, causing it to be retained.

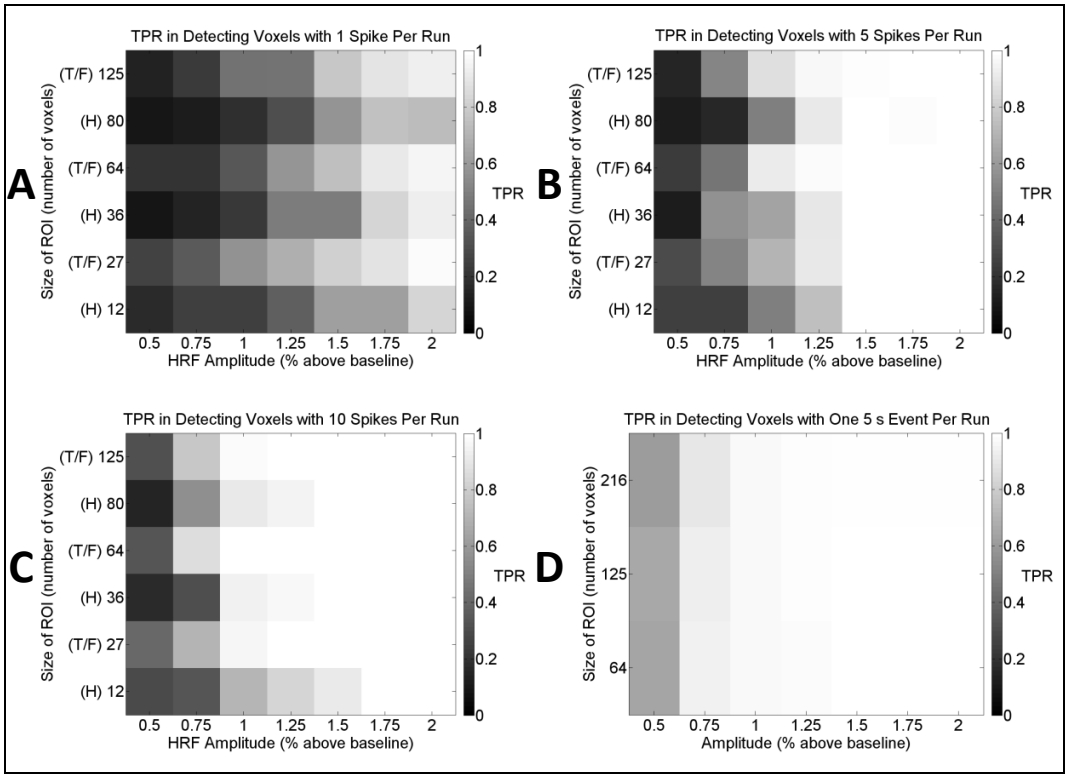

Figure 4.1-1 TPR in detecting simulated activity of different HRF amplitudes and ROI sizes in the case of (A) 1 spike per run, (B) 5 spikes per run, (C) 10 spikes per run, and (D) one 5 s event per run. As 2D-TCA creates a number of components for a given run, the TPR values shown are those associated with that component whose corresponding activation map best described the ROI of the simulated epileptic activity. Labels T, F, and H indicate brain areas containing the ROI (T = left temporal lobe, F = right frontal lobe, and H = left hippocampus); it should be reiterated that the 5 s event was only simulated in the right parietal lobe ROI.

Figure 4.1-2 shows the same data as in Figure 4.1-1 except collapsed across ROI sizes such that the TPR in detecting the various forms of simulated activity (i.e. 1 spike/run, 5 spikes/run, 10 spikes/run, one 5 s event) could be seen on the same plot. It was decided that for determining effective and consistent detection of a given form of simulated activity, 2D-TCA should produce an average TPR of at least 0.95. It can be seen from this plot that for effective detection by 2D-TCA, 1 spike per run is insufficient (maximum TPR of 0.897 for an HRF amplitude of 2%), while 5 spikes per run requires an HRF amplitude of at least 1.5% above baseline (corresponding TPR of 0.999), 10 spikes per run at least 1.25% (TPR of 0.959), and one 5 s event at least 1% (TPR of 0.976). Although 2D-TCA can consistently create components that precisely describe simulated epileptic activity with these

characteristics, it is important to keep in mind that it also creates many other components not associated with the simulated activity, some of whose corresponding t-maps may contain significant regions of activation. Figure 4.1-3 is a box plot showing, for runs simulated with an HRF amplitude of 1% or larger (i.e. the minimum HRF amplitude for which simulated activity was consistently detected), the number of components created by 2D-TCA whose corresponding activation maps contain, within the brain, clusters larger than a range of sizes. The number of activation maps steadily decreases with increasing cluster size thresholds until it starts to plateau around a cluster size threshold of 120 voxels. In fact, when applying a cluster size threshold of around 100 voxels or larger a median of only 2 or 3 components remain with significant activation. While this amounts to a significant reduction in the number of final components, even when applying a cluster size threshold of 220 voxels, which is larger than the largest simulated ROI (216 voxels), up to 3 components can still create significant activation that, in some cases, could be interpreted as arising from epileptic activity.

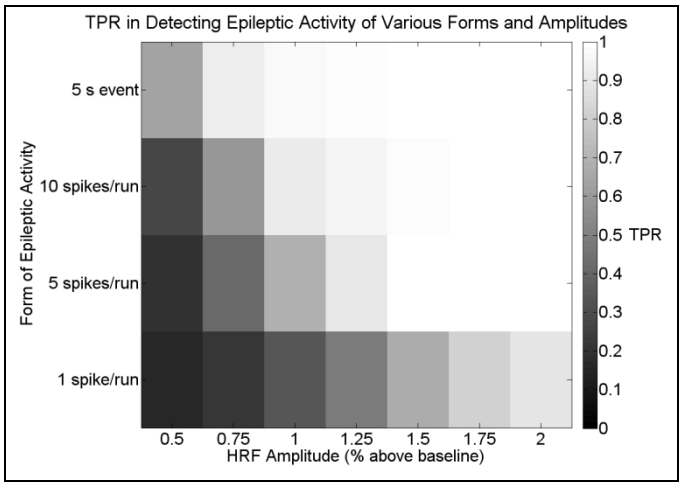

Figure 4.1-2 TPR for various HRF amplitudes and forms of epileptic activity. As 2D-TCA creates a number of components for a given run, the TPR values shown are those associated with that component whose corresponding activation map best described the ROI of the simulated epileptic activity.

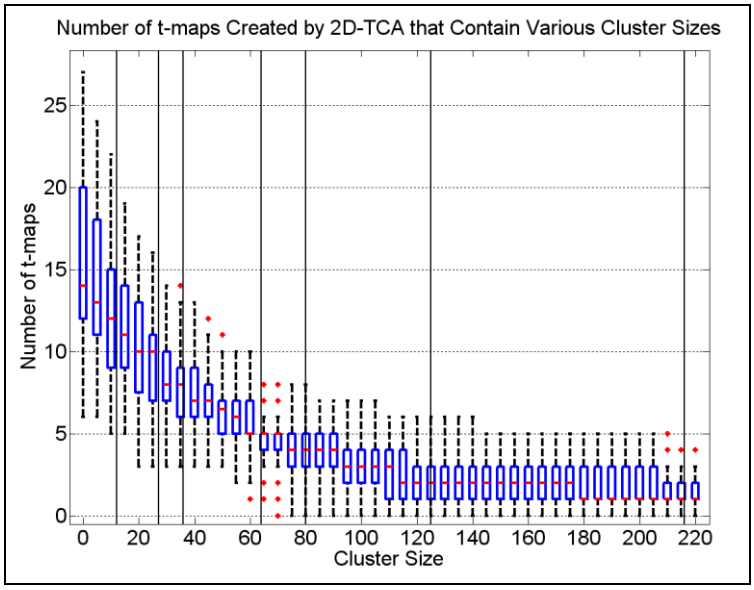

Figure 4.1-3 The number of components whose activation maps have cluster sizes above various threshold levels. Only tests using epileptic activity simulated with an HRF amplitude of 1% or more are considered. Vertical black lines indicate the simulated epileptic activity ROI sizes (i.e. 12, 27, 36, 64, 80, 125, and 216).

4.2 Performance in Detecting Epileptic Activity in Patients

Patient runs used to test the performance of 2D-TCA were selected based on the limits of detection found by the simulation tests. A run was selected if its EEG-fMRI results showed that there was clear activation and if the characteristics of this activation were similar to the consistently detectable forms of simulated activity described in section 4.1. However, to increase the number of analyzed patient runs, criteria, as found from the simulation results, were slightly relaxed (for example, very few patient runs would contain a single event lasting for about 5 s as was simulated). This meant that for a run to be selected at least 4 spikes or multiple prolonged events, whose total time of activity was 3 s or longer, needed to be detected in the EEG. In addition, the HRF amplitude associated with the activity, found from the EEG-fMRI results, needed to be at least around 1% above baseline. A total of 60 runs, 40 of which contained spikes and 20 of which contained prolonged events, were selected from across 20 patients. These 20 patients were selected from a larger pool of 43 patients (runs

recorded for the other 23 patients comprising this larger pool did contain activity which passed the above stated criteria).

Table 4.2-1 shows the performance of 2D-TCA, as compared to EEG-fMRI, in detecting spikes in patient data runs (each row represents a separate run) after cluster size thresholds of 15, 50, 100, and 200 have been applied to resulting t-maps (these thresholds were chosen based on the range of cluster sizes obtained from the EEG-fMRI results associated with spikes), while Table 4.2-2 shows the performance in detecting longer events for the same cluster size thresholds in addition to thresholds of 500 and 1000 (chosen based on the range of cluster sizes obtained from the EEG-fMRI results associated with long events); both tables also give the characteristics of the epileptic activity recorded in the run. A case in which 2D-TCA produced at least one t-map that passed the given cluster size threshold and, qualitatively, closely described what was seen in the EEG-fMRI results (i.e. similar regions of activation with spatially close peaks in activation), an example of which is shown in Figure 4.2-1, is indicated by a green cell, while a case in which no such similarity was seen for any of the t-maps is indicated by a red cell. A yellow cell indicates a case in which close similarity was not found, but that one of the t-maps created by 2D-TCA, that passed the given cluster size threshold, largely overlapped with what was seen in the EEG-fMRI results but did not include the maximum t-valued voxel (for example see Figure 4.2-2). Also given is the number of components created by 2D-TCA for each run, as well as the number whose corresponding t-maps actually contain significant activity within the brain, and the number that contain a cluster of activity larger than the a applied thresholds (note that performance before the 15 voxel cluster threshold is applied was not considered as any t-maps not passing this criteria could be considered as resulting from noise as t-maps are only created for those components who have a maximum voxel count of 15 or more). In terms of spike detection, Table 4.2-1 shows that in the majority of cases 2D-TCA is not able to

effectively detect what is found by EEG-fMRI, although it was found that there may be a slight dependence on the actual size of the activation (it was found that there was no dependency on number of spikes or HRF amplitude). On the other hand, by looking at Table 4.2-2, it seems that 2D-TCA is able to effectively detect what is found by EEG-fMRI when a run contains 2 or more interictal events of longer duration (on the order of a few seconds), while if only 1 event occurs it is essentially undetectable. For those runs in which 2D-TCA created a t-map that described what was seen by EEG-fMRI (i.e. what is indicated by a green cell), applying increasing cluster size thresholds generally didn't result in a deterioration of performance; this is also true for many of the runs for which 2D-TCA created a t-map whose activation overlapped what was seen by EEG-fMRI (i.e. what is indicated by a yellow cell). While performance was maintained for these runs, increasing the cluster size threshold reduced the number of t-maps that would be considered.

As mentioned, the output of 2D-TCA will be a set of t-maps, each of which hopefully describe a different form of activity. Figure 4.2-3(A) shows a box plot indicating the number of t-maps created for all 40 patient runs that contained spikes, as well as the number containing activity in the brain and the number containing a cluster of activity larger than the applied thresholds; Figure 4.2-3(B) shows the same data for all 20 patient runs that contained prolonged events. It is obvious that by increasing the cluster size threshold one reduces the number of t-maps that are considered. In an ideal case, by applying this cluster size threshold only a single t-map would remain, and that t-map would describe the epileptic activity. In fact, for 8 of the 15 runs in which similar activity was detected (indicated by green cells in Table 4.2-1 and Table 4.2-2), that map which closely described what was seen in the EEG-fMRI results contained the largest cluster of activity out of all 2D-TCA maps created for that run; however, in the 7 other cases this was not true and another form of activity showed larger

Epileptic Activity Characteristics (obtained from EEG-fMRI data)				2D-TCA Results					
Patient #	# of spikes	HRF Amplitude (% above baseline)	Size (voxels)	# of components	# of t-maps with activity	# of t-maps with cluster size of:			
						15	50	100	200
1	6	1.81	762	14	11	10	8	7	4
1	17	1.03	733	9	8	7	6	4	3
2	13	1.81	710	16	10	10	10	9	6
3	5	1.53	697	3	3	2	1	0	0
4	4	1.18	623	2	1	1	1	0	0
5	4	1.45	507	28	18	15	12	8	5
6	6	1.06	478	1	1	0	0	0	0
7	18	1.87	424	2	2	2	2	2	2
8	5	1.86	219	6	8	6	4	4	3
9	5	1.36	201	9	8	6	4	2	1
2	4	1.13	164	15	8	6	6	5	4
1	7	0.99	155	7	3	3	3	3	1
2	4	1.69	152	10	4	4	4	4	2
2	5	1.56	132	14	6	5	4	3	3
10	4	0.94	127	9	6	4	4	2	1
11	11	1.75	114	3	2	2	1	0	0
12	8	1.83	106	11	8	6	3	3	3
5	7	1.65	102	10	8	8	8	6	5
6	13	1.58	100	3	2	2	2	1	1
4	15	1.31	99	2	1	0	0	0	0
5	15	1.22	98	7	6	6	3	2	2
13	5	1.68	97	7	6	6	4	3	2
1	14	1.64	84	6	3	3	3	3	3
1	18	1.29	82	7	6	5	5	5	1
7	30	1.07	81	6	3	3	3	2	0
13	6	1.61	78	3	3	3	3	2	2
5	4	2.12	69	13	8	7	4	2	1
4	9	2.05	63	4	3	2	0	0	0
3	4	1.18	60	4	2	2	2	2	2
10	5	1.92	51	2	2	2	1	0	0
2	5	1.3	42	6	5	4	3	2	1
14	7	1.73	36	10	5	5	2	1	0
11	9	1.59	33	3	3	1	1	0	0
11	5	1.22	29	9	8	7	2	2	2
4	4	1.72	27	8	2	2	1	1	0
8	5	1.85	26	8	6	5	4	1	0
9	9	1.34	25	3	2	2	2	2	2
15	10	1.26	23	13	11	8	7	5	2
16	8	0.95	19	15	8	6	6	6	3
8	6	1.64	17	17	10	10	7	5	4

Table 4.2-1 Performance of 2D-TCA in detecting spikes in patient runs when different cluster size thresholds are applied. Runs are sorted by EEG-fMRI activation size from highest to lowest.

Epileptic Activity Characteristics (obtained from EEG-fMRI data)						2D-TCA Results							
Patient #	# of events	Average event duration (s)	Total time active (s)	HRF Amplitude (% above baseline)	Size (voxels)	# of components	# of t-maps with activity	# of t-maps with cluster size of:					
								15	50	100	200	500	1000
17	18	2.97	53.46	1.42	3341	10	7	6	6	6	4	1	0
17	9	3.19	28.71	1.33	3159	11	10	10	9	6	6	2	0
17	7	3.16	22.12	1.72	2845	13	13	8	6	3	2	2	0
7	4	2.65	10.6	1.14	1385	5	5	5	5	4	3	2	0
18	3	6.6	19.8	1.75	6469	5	3	2	2	2	2	0	0
19	3	2.63	7.89	1.36	3843	11	9	9	7	5	3	0	0
7	3	2.27	6.81	1.23	1312	18	6	5	3	3	2	1	1
17	2	6.3	12.6	3.01	3478	8	8	7	5	4	4	1	0
20	2	2.57	5.14	1.95	726	14	11	10	6	4	3	2	1
20	2	2.38	4.76	1.08	621	13	13	8	3	3	1	0	0
19	2	4.7	9.4	1.58	3678	10	8	7	6	5	4	1	0
19	2	3.8	7.6	1.28	3556	6	3	3	3	2	2	1	0
17	1	3.1	3.1	1.29	2142	6	6	5	3	3	2	0	0
1	1	4.2	4.2	2.41	71	8	8	7	5	4	4	3	1
20	1	4.07	4.07	1.06	597	3	3	3	2	2	2	1	0
7	1	6.2	6.2	1.22	1440	5	5	2	2	1	0	0	0
17	1	5	5	1.12	2798	7	6	5	2	2	1	1	0
19	1	4.5	4.5	1.07	3802	4	2	2	1	1	1	0	0
7	1	3.4	3.4	1.45	1189	16	13	10	4	3	3	0	0
7	1	3.5	3.5	1.29	845	16	10	10	5	2	2	0	0

Table 4.2-2 Performance of 2D-TCA in detecting longer interictal events when different cluster size thresholds are applied. Runs are sorted by number of events detected in the EEG from highest to lowest.

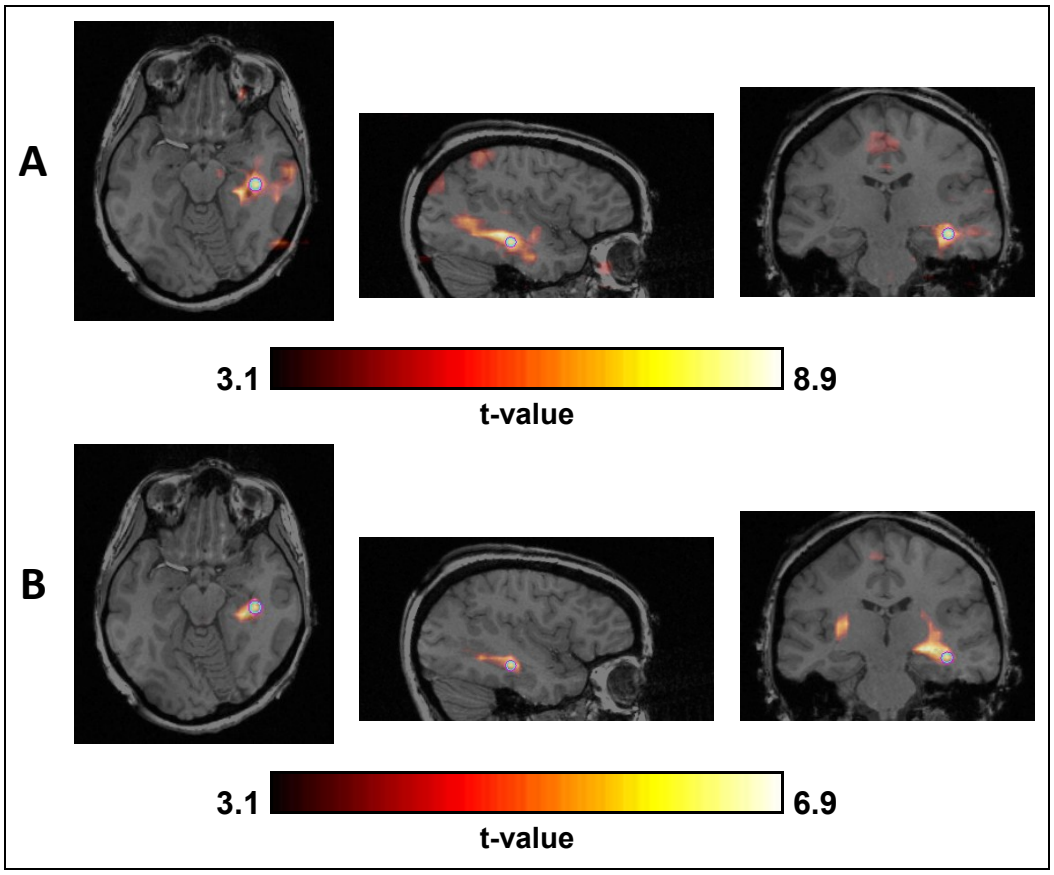

Figure 4.2-1 Example of strong concordance (indicated by green in Table 4.2-1 and Table 4.2-2) between the t-map created by (A) EEG-fMRI and (B) one of the 2D-TCA components. The purple circle indicates the exact voxel whose slices are shown.

activation than the detected epileptic activity. Looking at Figure 4.2-3(A) it can be seen that, for all runs that contained spikes, even when a cluster size threshold of 200 voxels is applied (which only provided a median of 2 remaining t-maps), up to 6 t-maps can remain. Likewise, for all runs that contained prolonged events, when a cluster size threshold of 500 voxels is applied (which provided a median of 1 remaining t-map), up to 3 t-maps can remain.

While some of these components may include those that describe epileptic activity, many of them may describe other forms of activity that, without knowledge of EEG-fMRI results, may or may not be interpreted as arising from epileptic activity depending on the region of activation and the other information available (e.g. EEG, clinical observations, anatomical MRI, etc.).

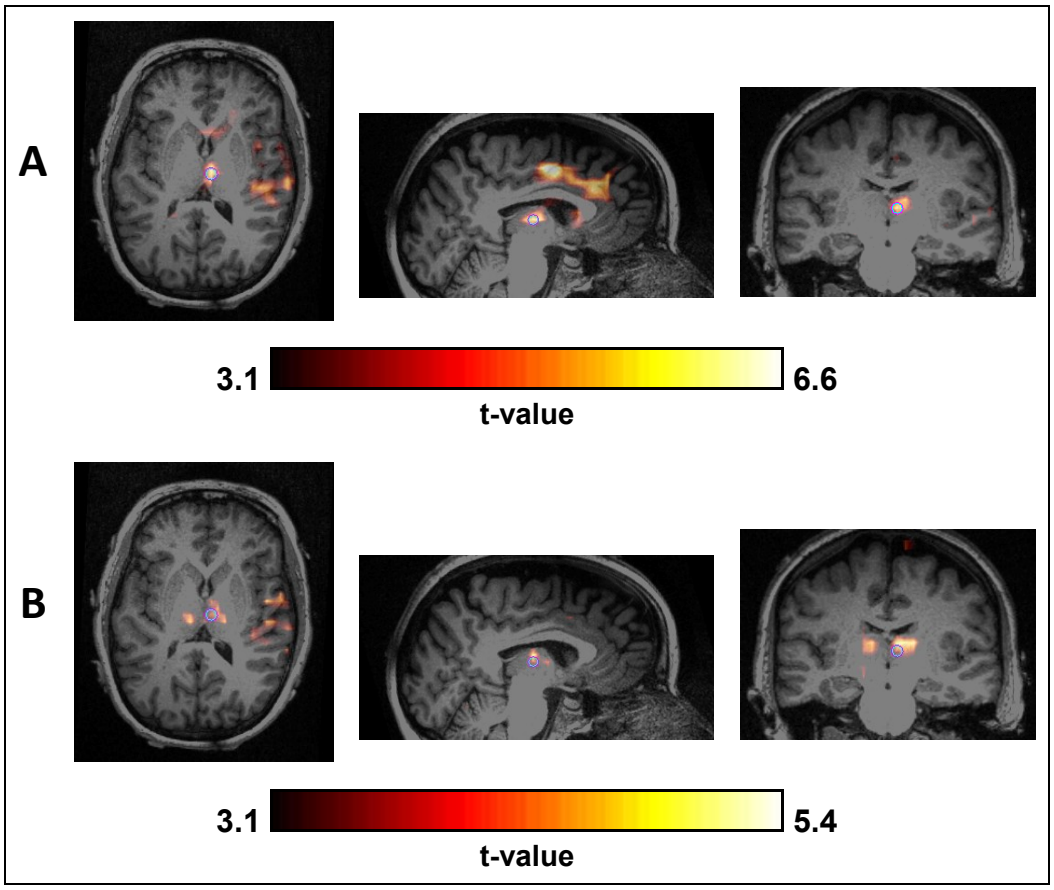

Figure 4.2-2 Example of some overlap (indicated by yellow in Table 4.2-1 and Table 4.2-2) between the t-map created by (A) EEG-fMRI and (B) one of the 2D-TCA components. The purple circle indicates the exact voxel whose slices are shown. Note the activity in the frontal lobe in the sagittal slice of the EEG-fMRI map and its absence in the 2D-TCA map.

Among these extra components, when a cluster size threshold of 15 voxels was applied, two were found to be common across a number of runs from different patients. The first, which was found to occur in 14 of the 60 runs, consisted of activation inside or close to the ventricles, an example of which is shown in Figure 4.2-4. It is believed that these maps may describe BOLD signal changes arising from residual movement artifacts, the patient's respiration, or the patient's heartbeat as the ventricles are susceptible to these forms of noise (see section 2.3.6). The second form of common activity, which was found to occur in 12 of the runs, consisted of bilateral activity close, often slightly posterior, to the central sulcus, an example of which is shown in Figure 4.2-5. It is believed that these maps may describe motor or somatosensory activity that occurred in the

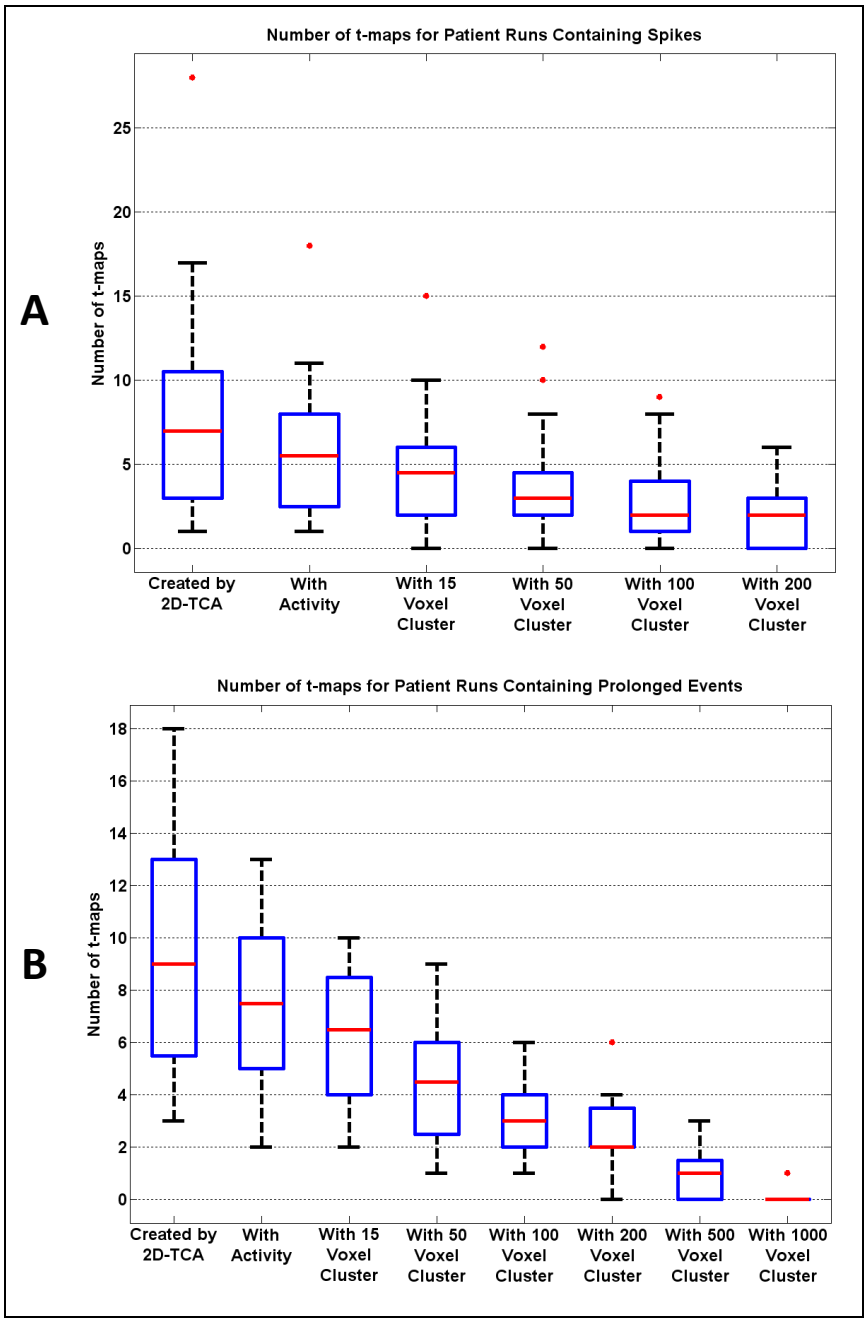

Figure 4.2-3 Box plot showing, for (A) all 40 patient runs containing spikes and (B) all 20 patient runs containing prolonged events, the number of t-maps created, number of those created that contain activity in the brain, and number of those that contain a cluster of activity larger than the applied threshold levels.

patient, or was imagined by them, during the run. In fact, this region of activation is very similar to the common motor cortex activation observed in (Morgan et al 2008). In that study they state that the cause of the activity is not

completely understood, but that it is most likely due to some sort of motor activity by the subject even though no such deliberate task was required of them.

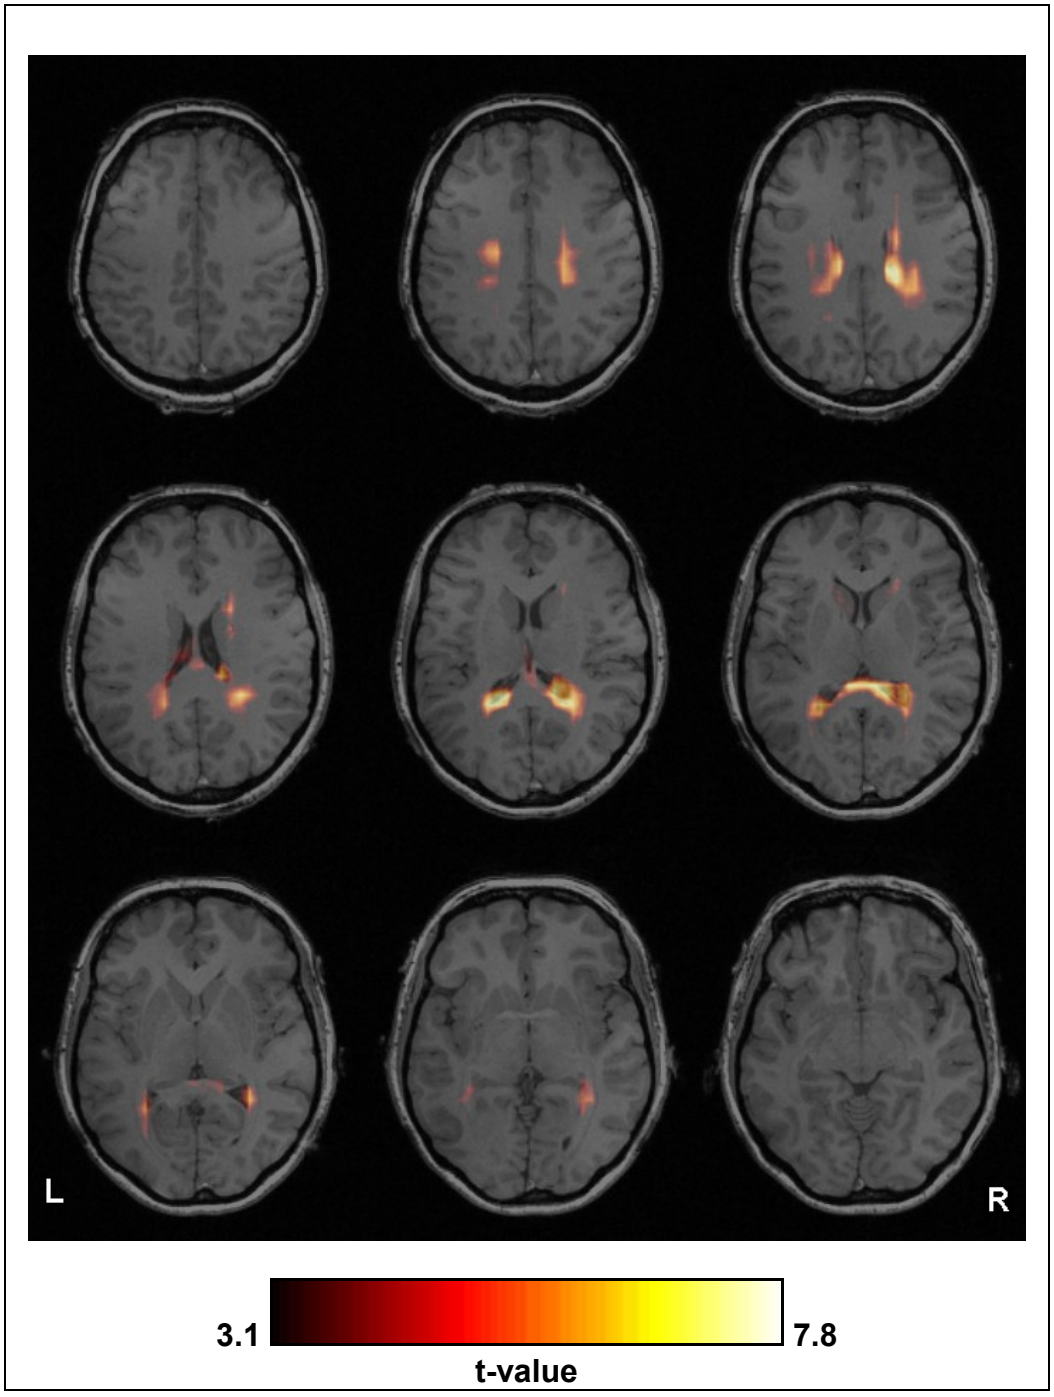

Figure 4.2-4 Example of commonly found activation in the ventricles.

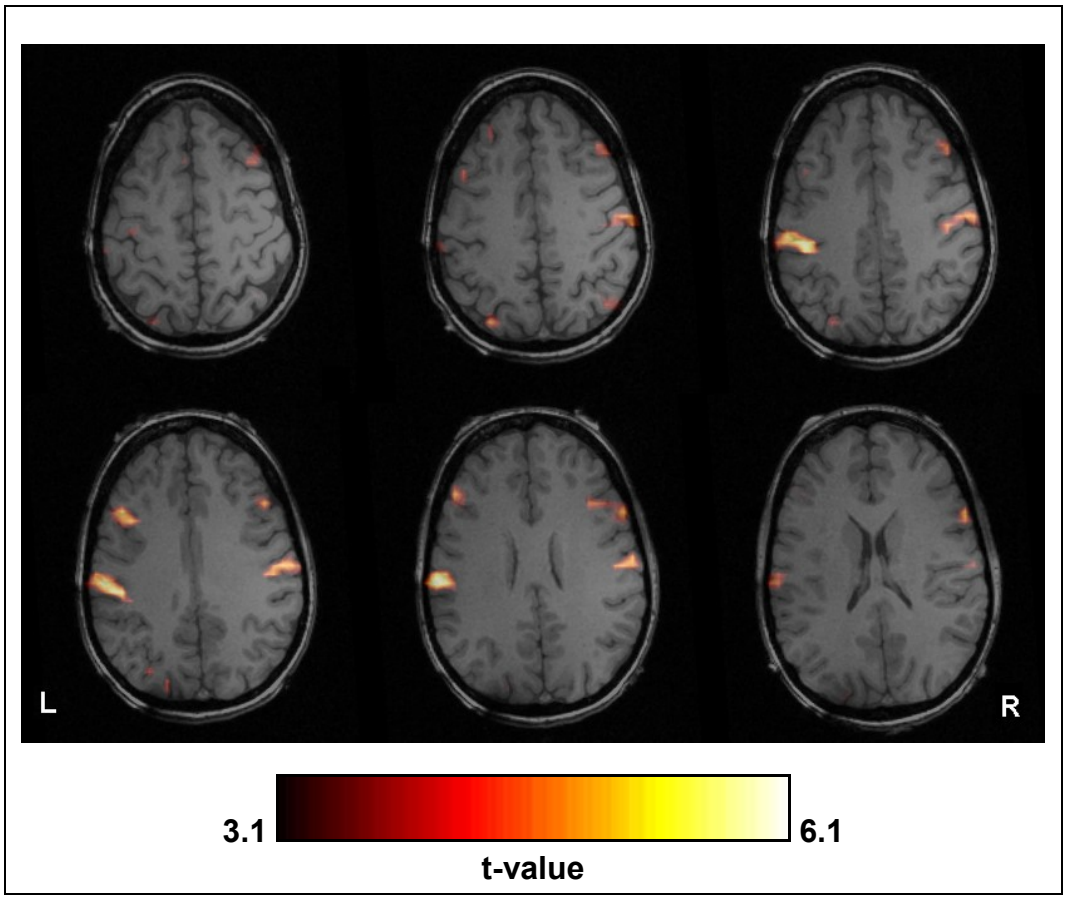

Figure 4.2-5 Example of commonly found activation near (in this case slightly posterior to) the central sulcus.

4.2.1 Analysis by Neurologist

Runs for which 2D-TCA created a t-map that closely described what was seen in the EEG-fMRI results (i.e. those 16 rows with green in Table 4.2-1 and Table 4.2-2; only those t-maps created for these runs that had a cluster size of at least 15 voxels were considered) were given to a neurologist for further analysis. This was done to determine, without the knowledge of EEG-fMRI results, the number of t-maps created by 2D-TCA that may be interpreted as describing epileptic activity within a given patient and the number that can be ignored on the basis that they show regions of activation that are highly unlikely to be generated by a given patient's epileptic activity. Figure 4.2-6 shows a box plot indicating the number of t-maps created in these cases, as well as the number containing activity in the brain, the number containing a cluster of activity of at least 15

voxels, and the number that the neurologist indicated may be interpreted as arising from potential epileptic activity. While for all these runs it was known that only one of the t-maps created by 2D-TCA closely described what was seen with EEG-fMRI, the neurologist found that a median of 2 and a maximum of 5, each describing a different region of activation, could not be ignored and therefore could be considered as potentially being associated with the epileptic activity. For five of these runs the neurologist gave a score of 4 or 5 (values given for a t-map which describes areas of activation that are most likely resulting from epileptic activity) to a 2D-TCA created t-map that showed similar activation to what was seen in the EEG-fMRI results, while in three other runs the neurologist gave a score of 4 to a t-map that was known not to show similar activity to what was obtained by EEG-fMRI; the neurologist gave a maximum score of 3 (value given for a t-map which describes areas of activation that may or may not be resulting from epileptic activity) for t-maps created for the remaining 8 runs. So, while for some runs it was possible to easily select those 2D-TCA created activation maps that describe epileptic activity, for 11 of the 16 analyzed runs this was not possible.

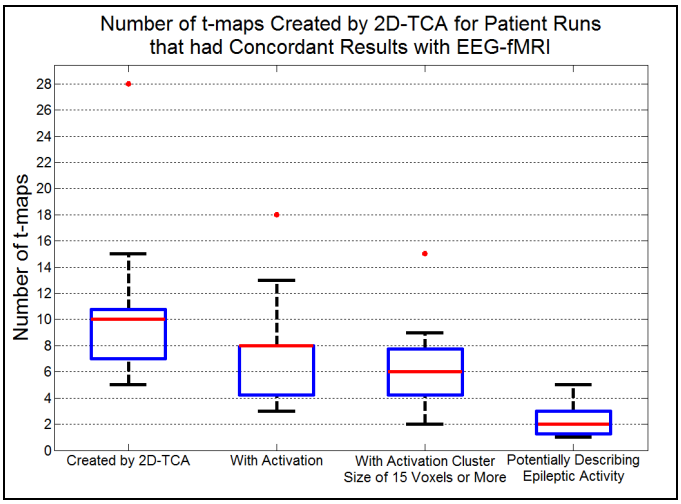

Figure 4.2-6 Box plot showing, for those runs that 2D-TCA created a t-map that closely described what was seen in the EEG-fMRI results, number of t-maps created, number of those created that contain significant activity in the brain, number of those with a cluster of activity of 15 voxels or larger, and number that, in the opinion of the neurologist, could not be ignored and therefore may describe epileptic activity.

5 Discussion

The main advantage to localizing epileptic activity in fMRI data, for example by EEG-fMRI, is the ability to do so throughout the entire brain with high spatial precision. However, if this can be done without dependence on events detected in a simultaneously recorded EEG, the scanning process would be much more manageable and activity restricted to deep brain structures, which is undetectable by EEG, may actually be detected. This study investigated the ability of 2D-TCA, a data based method that detects transient activity in fMRI data independently of EEG, to detect BOLD signal changes resulting from epileptic activity of various forms in both simulated and patient data. Activity characteristics that were investigated included event frequency/duration, spatial extent, and associated HRF amplitude. When applied to the simulated data it was found that 2D-TCA functioned similarly in detecting activity of various spatial extents, but that it could only consistently detect activity in a run containing 5 spikes/run, with an associated HRF amplitude of 1.5% above baseline or larger, 10 spikes/run, with an HRF amplitude of 1.25% or larger, or one 5 s long event, with an HRF of 1% or larger. In all other cases, for example if there were fewer than 5 spikes/run or if the HRF amplitude was too small, 2D-TCA could not effectively detect the activity. When applied to patient data that contained activity similar to what was found to be detectable in the simulated data, it was found that only epileptic activity for which the EEG showed multiple prolonged events (each on the order of a few seconds) could be consistently detected by 2D-TCA, while the detection of activity consisting of interictal spikes, although in some cases very accurate, was generally inconsistent. This better performance in detecting longer events may not only be due to the longer duration of activation (which would allow for a larger number of detected events), but also because for the same HRF, a prolonged event will cause a larger change in the BOLD signal than a spike will. It is also interesting to note that t-maps created by 2D-TCA,

which showed very similar activation to what was obtained by EEG-fMRI, often had more focal activation regions than their EEG-fMRI counterparts. This may indicate that 2D-TCA is more sensitive to the initial source of epileptic activity than EEG-fMRI is, a quality that would require further investigation to confirm. It should be noted that when creating t-maps from multiple regressors (i.e. components) the common practice is to pass all regressors through the same GLM analysis. However, in this study, as the number of components created by 2D-TCA (greater than 5 and up to 15) was generally larger than the number used in typical event related fMRI studies, t-maps were created by passing each regressor through a separate GLM analysis. This was done to insure that any possible co-linearity between components did not diminish the level of activation seen in a given component's t-map. The downside of this is that the model's accuracy in describing activity may suffer. For example, if one voxel is actually affected by two separate types of activity, each described by a separate regressor, then the activity in the voxel is a function of both regressors and the model will be more accurate (i.e. produce a smaller error) if both regressors are passed through the same GLM, as opposed to passing each separately.

The issue then arises: if 2D-TCA will only detect epileptic activity that has the above described characteristics which, in the case of a patient, are uncontrollable, how can it be determined, in the absence of EEG, if activity is expected to be detected in a given patient? In fact, given these conditions, it would not be possible to determine beforehand if a patient's epileptic activity is expected to be detected with 2D-TCA. Rather, the above described characteristics, which concern the HRF amplitude and frequency/length of the activity, might be incorporated into the analysis of the t-maps generated by 2D-TCA. For example, the HRF amplitude associated with activity described in each t-map could be calculated by taking the average peak amplitude of all voxels that are clustered to the same final component. Those components for which the HRF

amplitude is not deemed as sufficient could then be ignored. To determine whether or not activity described by a given t-map consisted of frequent and/or long enough activity one would simply have to look at the actual time course of the final component used to create that t-map. If the component's time course does not describe frequent or long enough activity, its corresponding t-map could then potentially be ignored.

Another possible method to reduce the number of considered t-maps would be to apply a cluster size threshold. For cases in which 2D-TCA created a t-map that closely described what was seen by EEG-fMRI, applying larger and larger cluster size thresholds led to fewer and fewer t-maps, but retained that which described the epileptic activity. Unfortunately, as the size of the activation region associated with a patient's epileptic activity is not known before hand, it is hard to determine an optimal cluster size threshold that would maximally reduce the number of t-maps while retaining that which describes the epileptic activity. It would be inappropriate to exclusively select that t-map which showed the largest activation region as it would not necessarily be the one that describes the epileptic activity. The best way to treat the situation may be to reduce the number of considered t-maps, to something reasonable for visual inspection, by applying a cluster size threshold of 15 (for those runs in which activity similar to the EEG-fMRI results was obtained, applying this threshold resulted in a median of 6 t-maps). This reduced set could then be investigated by a neurologist to determine whether any of the included t-maps show activity of interest. In this study it was found that while a neurologist is able to ignore some of these maps on the basis that they describe activation regions that are highly unlikely to be associated with the syndrome of epilepsy seen in a given patient, in most cases more than one will remain, leading to an uncertainty as to which of the remaining activation maps describes the activity of interest. In addition, it was found that for 5 out of the 16 runs for which 2D-TCA created a component which

closely described what was seen by EEG-fMRI, the neurologist, blind to the EEG-fMRI results, could not confidently qualify the t-map as describing regions of activation associated with the given patient's epilepsy. This lack of confidence in determining which t-maps are and are not associated with epilepsy results from 2D-TCA's lack of specificity in detecting activity associated with epilepsy (EEG-fMRI achieves this specificity by finding BOLD correlates to epileptic activity detected in the patient's EEG). As such, 2D-TCA can only be effectively used to validate localization by other means or to create hypotheses as to where epileptic activity may be occurring.

5.3 Advantages Compared to EEG-fMRI

As mentioned, one of the major advantages of 2D-TCA when it is compared to EEG-fMRI is its potential to detect activity restricted to deep brain structures. While this was evident when 2D-TCA was applied to the simulated data (i.e. detection of activity in the right hippocampus ROI), it was not as clear from the patient data. This study saw three cases in which 2D-TCA applied to patient data created a t-map that the neurologist decided most likely described a region associated with epileptic activity but was not similar to what was seen in the EEG-fMRI produced map. However, it is expected that the areas of activation in all three cases would have been detected in the EEG as all were relatively superficial and consisted of an activated area of cortex larger than the minimum 10 cm^2 required for EEG detection (Tao et al 2007) (why similar regions were not detected by EEG-fMRI may be a result of inadequate EEG markings, or, more likely, indicative that these regions were in fact not associated with epileptic activity). In fact, depth electrode recordings, which confirm that a patient's epileptic activity is restricted to deep brain structures, would most likely be required for comparison to determine whether 2D-TCA is able to detect such activity or not; unfortunately, the patients included in this study had no such recordings.

Another point to consider is the assumption that EEG-fMRI requires one to impose in terms of the delay of the HRF, something that can be different from one person to another or even from one brain structure to another. Some techniques try to overcome this problem by processing the same fMRI data multiple times, each time with an HRF of different delay. On the other hand, 2D-TCA requires no such assumption as it is a data based technique that detects events straight from the BOLD signal itself; only the exact shape of the HRF must be assumed, something that does not show nearly as much variability across different people or brain structures.

5.4 Comparison to ICA

Although independent component analysis (ICA) was not carried out in this study, the performance of 2D-TCA in detecting transient BOLD activity as compared to ICA has been previously investigated by Morgan et al. (Morgan et al 2008). They found that detection by 2D-TCA was comparable to that by ICA (2D-TCA provided slightly worse sensitivity), but with significantly fewer detected components (on the order of 5 compared to 10 or 100) and hence a higher specificity. Based on the 2D-TCA results obtained in this study and the findings of Morgan et al., we would predict a similar outcome; that is, applying ICA to the data used in this study would in fact provide better sensitivity than 2D-TCA did, but the number of extra components it creates compared to the number created by 2D-TCA makes 2D-TCA the more practical technique of the two.

5.5 Future of TCA for Detection of Epileptic Activity

Although it may be considered that 2D-TCA produces less certain results than EEG-fMRI, as the t-maps it creates lack specificity to epileptic activity, some changes not considered in this study may provide improvements in the

technique's ability to detect epileptic activity, particularly of lower amplitude and event frequency.

The first change involves the method by which event time courses are clustered into the 2D histogram. By the current method, event time courses are clustered based on the time point at which the voxel achieves its maximum value. An issue which arises is that if a voxel which is associated with epileptic activity is also associated with some other sort of larger activation during the scanning run, this other activation will define when its peak value occurs, and therefore how its event time course is clustered into the 2D histogram. This issue will reduce the ability of the algorithm to create final components that correlate with specific forms of activity. One idea to overcome this problem would be to cluster event time courses based on how similar they are to one another (e.g. by correlation or k-means cluster analysis) rather than by when their associated voxels reach their maximum values. In addition, spatial constraints, based on proximity and/or anatomy, could be imposed to define which event time courses can and cannot be clustered together. By implementing such a technique, which treats the temporal clustering step much more rigorously, it is believed that the final components created by 2D-TCA will more accurately describe different forms of activity seen in a given run.

Another idea to consider would be to investigate the ability of 2D-TCA to detect epileptic activity within a restricted ROI. As applied in this study, 2D-TCA processed all voxels within the brain. However, its performance may improve if it is only applied to a specific region of the brain (for example one hemisphere, or lobe) within which the epileptogenic zone is expected to exist (i.e. based on what is observed from a patient's EEG and other evaluation techniques). As fewer voxels would be processed, fewer voxels would be added to a given row in the 2D histogram simply by chance, and therefore, 2D-TCA would create final

components that more accurately describe the detected activity. In addition, as activity occurring outside of the specified ROI would be ignored, it is expected that the number of t-maps would be reduced. It should be noted that in such a situation a new, more appropriate, insignificant component threshold value may need to be applied due to the fewer number of voxels being processed.

Finally, the ability of 2D-TCA to detect epileptic activity across multiple fMRI runs should be investigated. Many EEG-fMRI techniques increase their statistical strength by combining the results obtained from separate runs recorded during the same scanning session. This idea could be incorporated into 2D-TCA in two ways. The first would be to simply concatenate multiple runs into a single scan consisting of voxels with much longer time courses. 2D-TCA could then be applied onto this single scan, although some steps in the 2D-TCA algorithm would have to be appropriately modified to account for the increased number of time samples, the differences in baseline that may exist from one run to another, and other issues of discontinuity that may arise when concatenating runs together. This would potentially help detect activity consisting of fewer events per run as more events would be able to be detected over this longer period. The second method would be to apply 2D-TCA on a run by run basis, but to then compare t-maps created for separate runs from the same scan session. While this would not improve 2D-TCA's ability to detect activity consisting of few events, it may aid in determining which t-maps describe regions of activation associated with epileptic activity and which do not. It is expected that while t-maps that describe the epileptogenic zone may be created for a number of runs, other t-maps that describe other transient activity might be specific to one run. As such, those t-maps for which a similar region of activation is seen across many runs may be retained, while those for which such similarity is not seen may be ignored. This would only work for patients who consistently have a frequent number of events such that many would occur during each 6 min run.

6 Summary and Conclusion

This thesis has provided information related to the localization of epileptic activity within patients, specifically investigating the use of a relatively new technique, 2D-TCA, for this purpose. First, information was provided regarding the condition of epilepsy, specifically looking at common syndromes of epilepsy, as well as current methods of treatment. This was then followed by a review of the most common epileptic activity recording modality, EEG, which included information on the physical basis of its acquisition, as well as how it can be used to detect and localize epileptic activity. A review of fMRI was then presented, specifically looking at the physical basis of how the MR signal is recorded, the hemodynamic changes in the brain caused by neuronal activity, how these changes can be recorded through the BOLD signal, the most common fMRI data analysis method, the GLM, and the sources that cause noise in a typical fMRI scan. This was then followed by a description of common recording and analysis methods for EEG-fMRI, a technique used to localize epileptic activity with high resolution, as well as the advantages and disadvantages associated with using this technique. The initial development of TCA, a relatively new technique implemented to detect transient activity in fMRI data, was then presented along with information concerning its first applications to detect epileptic activity independently of the EEG, as well as the development of 2D-TCA for this same purpose.

This thesis then provided the methods, results, and discussion associated with a study carried out to investigate the ability of 2D-TCA to effectively detect epileptic activity of various forms in both simulated and patient fMRI data. It was found that in certain circumstances, namely if enough events were recorded during an fMRI run and the HRF associated with those events was of adequate amplitude, 2D-TCA could consistently detect epileptic activity. However, due to

its lack of specificity in creating activation maps that describe epileptic activity, it was concluded that 2D-TCA can only be effectively used to validate localization of epileptic activity by other means or to create hypotheses as to where this activity may be occurring. This being said, the ability to confidently detect epileptic activity with 2D-TCA should not be necessarily ruled out as it may be further developed through the application of certain modifications that could potentially improve its performance, some examples of which have been presented in this thesis.

References

Aguirre GK, Zarahn E, D'Esposito M. 1998. The variability of human, BOLD hemodynamic responses. *Neuroimage* 8:360-9

Allen PJ, Josephs O, Turner R. 2000. A method for removing imaging artifact from continuous EEG recorded during functional MRI. *Neuroimage* 12:230-9

Aurlien H, Gjerde IO, Aarseth JH, Eldoen G, Karlsen B, et al. 2004. EEG background activity described by a large computerized database. *Clin Neurophysiol* 115:665-73

Baker GA, Jacoby A, Buck D, Stalgis C, Monnet D. 1997. Quality of life of people with epilepsy: a European study. *Epilepsia* 38:353-62

Bandettini PA, Kwong KK, Davis TL, Tootell RB, Wong EC, et al. 1997. Characterization of cerebral blood oxygenation and flow changes during prolonged brain activation. *Hum Brain Mapp* 5:93-109

Behzadi Y, Liu TT. 2005. An arteriolar compliance model of the cerebral blood flow response to neural stimulus. *Neuroimage* 25:1100-11

Benar C, Aghakhani Y, Wang Y, Izenberg A, Al-Asmi A, et al. 2003. Quality of EEG in simultaneous EEG-fMRI for epilepsy. *Clin Neurophysiol* 114:569-80

Berg AT, Langfitt JT, Spencer SS, Vickrey BG. 2007. Stopping antiepileptic drugs after epilepsy surgery: a survey of U.S. epilepsy center neurologists. *Epilepsy Behav* 10:219-22

Birbeck GL, Hays RD, Cui X, Vickrey BG. 2002. Seizure reduction and quality of life improvements in people with epilepsy. *Epilepsia* 43:535-8

Blume WT. 2003. Diagnosis and management of epilepsy. *CMAJ* 168:441-8

Bourgeois BFD. 2001. General concepts of medical intractability. In *Epilepsy surgery*, ed. H Lüders, YG Comair, pp. 63-8. Philadelphia: Lippincott Williams & Wilkins

Boynton GM, Engel SA, Glover GH, Heeger DJ. 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16:4207-21

Buxton RB. 2002. *Introduction to functional magnetic resonance imaging : principles and techniques*. Cambridge, U.K. ; New York: Cambridge University Press. xi, 523 p. pp.

Chen JJ, Pike GB. 2009. Origins of the BOLD post-stimulus undershoot. *Neuroimage* 46:559-68

Dagli MS, Ingeholm JE, Haxby JV. 1999. Localization of cardiac-induced signal change in fMRI. *Neuroimage* 9:407-15

Engel J. 1993. *Surgical treatment of the epilepsies*. New York: Raven Press. xxix, 786 p. pp.

Engel J, Jr. 1984. A practical guide for routine EEG studies in epilepsy. *J Clin Neurophysiol* 1:109-42

Engel J, Jr. 2006. Report of the ILAE classification core group. *Epilepsia* 47:1558-68

Frahm J, Kruger G, Merboldt KD, Kleinschmidt A. 1996. Dynamic uncoupling and recoupling of perfusion and oxidative metabolism during focal brain activation in man. *Magn Reson Med* 35:143-8

Fransson P, Kruger G, Merboldt KD, Frahm J. 1999. MRI of functional deactivation: temporal and spatial characteristics of oxygenation-sensitive responses in human visual cortex. *Neuroimage* 9:611-8

Friston KJ, Jezzard P, Turner R. 1994. Analysis of functional MRI time-series. *Human Brain Mapping* 1:153-71

Frucht MM, Quigg M, Schwaner C, Fountain NB. 2000. Distribution of seizure precipitants among epilepsy syndromes. *Epilepsia* 41:1534-9

Gao JH, Yee SH. 2003. Iterative temporal clustering analysis for the detection of multiple response peaks in fMRI. *Magn Reson Imaging* 21:51-3

Genovese CR, Lazar NA, Nichols T. 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage* 15:870-8

Glover GH. 1999. Deconvolution of impulse response in event-related BOLD fMRI. *Neuroimage* 9:416-29

Glover GH, Li TQ, Ress D. 2000. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magn Reson Med* 44:162-7

Gotman J, Grova C, Bagshaw A, Kobayashi E, Aghakhani Y, Dubeau F. 2005. Generalized epileptic discharges show thalamocortical activation and

suspension of the default state of the brain. *Proc Natl Acad Sci U S A* 102:15236-40

Gotman J, Kobayashi E, Bagshaw AP, Benar CG, Dubeau F. 2006. Combining EEG and fMRI: a multimodal tool for epilepsy research. *J Magn Reson Imaging* 23:906-20

Gregory RP, Oates T, Merry RT. 1993. Electroencephalogram epileptiform abnormalities in candidates for aircrew training. *Electroencephalogr Clin Neurophysiol* 86:75-7

Grova C, Daunizeau J, Lina JM, Benar CG, Benali H, Gotman J. 2006. Evaluation of EEG localization methods using realistic simulations of interictal spikes. *Neuroimage* 29:734-53

Hamandi K, Salek Haddadi A, Liston A, Laufs H, Fish DR, Lemieux L. 2005. fMRI temporal clustering analysis in patients with frequent interictal epileptiform discharges: comparison with EEG-driven analysis. *Neuroimage* 26:309-16

Helmholtz H. 1853. Ueber einige Gesetze der Vertheilung elektrischer Ströme in körperlichen Leitern, mit Anwendung auf die thierisch-elektrischen Versuche (Schluss.). *Annalen der Physik und Chemie* 165:353-77

Henson RN, Buchel C, Josephs O, Friston KJ. 1999. The slice-timing problem in event-related fMRI. *Neuroimage* 9:125

Huettel SA, Song AW, McCarthy G. 2004. *Functional magnetic resonance imaging*. http://www.loc.gov/catdir/toc/ecip0413/2004000500.html

Ives JR, Warach S, Schmitt F, Edelman RR, Schomer DL. 1993. Monitoring the patient's EEG during echo planar MRI. *Electroencephalogr Clin Neurophysiol* 87:417-20

Krakow K, Allen PJ, Symms MR, Lemieux L, Josephs O, Fish DR. 2000. EEG recording during fMRI experiments: image quality. *Hum Brain Mapp* 10:10-5

Kruger G, Glover GH. 2001. Physiological noise in oxygenation-sensitive magnetic resonance imaging. *Magn Reson Med* 46:631-7

Kwan P, Brodie MJ. 2000. Early identification of refractory epilepsy. *N Engl J Med* 342:314-9

Kybic J, Clerc M, Faugeras O, Keriven R, Papadopoulo T. 2006. Generalized head models for MEG/EEG: boundary element method beyond nested volumes. *Phys Med Biol* 51:1333-46

Lange N, Zeger SL. 1997. Non-Linear Fourier Time Series Analysis for Human Brain Mapping by Functional Magnetic Resonance Imaging. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 46:1-29

Lazeyras F, Zimine I, Blanke O, Perrig SH, Seeck M. 2001. Functional MRI with simultaneous EEG recording: feasibility and application to motor and visual activation. *J Magn Reson Imaging* 13:943-8

Lemieux L, Allen PJ, Franconi F, Symms MR, Fish DR. 1997. Recording of EEG during fMRI experiments: patient safety. *Magn Reson Med* 38:943-52

Liu Y, Gao JH, Liotti M, Pu Y, Fox PT. 1999. Temporal dissociation of parallel processing in the human subcortical outputs. *Nature* 400:364-7

Liu Y, Gao JH, Liu HL, Fox PT. 2000. The temporal response of the brain after eating revealed by functional MRI. *Nature* 405:1058-62

Lund TE, Larsson HBW. 1999. Spatial distribution of low-frequency noise in fMRI. *Proceedings of the 7th Annual Meeting of ISMRM*:1705

Lund TE, Madsen KH, Sidaros K, Luo WL, Nichols TE. 2006. Non-white noise in fMRI: does modelling have an impact? *Neuroimage* 29:54-66

Malmivuo J, Plonsey R. 1995. *Bioelectromagnetism : principles and applications of bioelectric and biomagnetic fields*. New York: Oxford University Press. xxii, 482 p. pp.

Mandeville JB, Marota JJ, Ayata C, Zaharchuk G, Moskowitz MA, et al. 1999. Evidence of a cerebrovascular postarteriole windkessel with delayed compliance. *J Cereb Blood Flow Metab* 19:679-89

Marsan CA, Zivin LS. 1970. Factors related to the occurrence of typical paroxysmal abnormalities in the EEG records of epileptic patients. *Epilepsia* 11:361-81

Merlet I, Gotman J. 1999. Reliability of dipole models of epileptic spikes. *Clin Neurophysiol* 110:1013-28

Morgan VL, Gore JC. 2009. Detection of irregular, transient fMRI activity in normal controls using 2dTCA: comparison to event-related analysis using known timing. *Hum Brain Mapp* 30:3393-405

Morgan VL, Gore JC, Abou-Khalil B. 2010. Functional epileptic network in left mesial temporal lobe epilepsy detected using resting fMRI. *Epilepsy Res* 88:168-78

Morgan VL, Li Y, Abou-Khalil B, Gore JC. 2008. Development of 2dTCA for the detection of irregular, transient BOLD activity. *Hum Brain Mapp* 29:57-69

Morgan VL, Price RR, Arain A, Modur P, Abou-Khalil B. 2004. Resting functional MRI with temporal clustering analysis for localization of epileptic activity without EEG. *Neuroimage* 21:473-81

Mosher JC, Leahy RM, Lewis PS. 1999. EEG and MEG: forward solutions for inverse methods. *IEEE Trans Biomed Eng* 46:245-59

Noachtar S, Binnie C, Ebersole J, Mauguiere F, Sakamoto A, Westmoreland B. 1999. A glossary of terms most commonly used by clinical electroencephalographers and proposal for the report form for the EEG findings. The International Federation of Clinical Neurophysiology. *Electroencephalogr Clin Neurophysiol Suppl* 52:21-41

Noachtar S, Borggraefe I. 2009. Epilepsy surgery: a critical review. *Epilepsy Behav* 15:66-72

Noachtar S, Remi J. 2009. The role of EEG in epilepsy: a critical review. *Epilepsy Behav* 15:22-33

Noachtar S, Winkler PA, Luders HO. 2003. Neurological disorders : course and treatment. In *Surgical therapy of epilepsy*, pp. 235-44. Amsterdam: Academic Press

Ogawa S, Lee TM, Nayak AS, Glynn P. 1990. Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. *Magn Reson Med* 14:68-78

Ogawa S, Menon RS, Tank DW, Kim SG, Merkle H, et al. 1993. Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model. *Biophys J* 64:803-12

Pauling L, Coryell CD. 1936. The Magnetic Properties and Structure of Hemoglobin, Oxyhemoglobin and Carbonmonoxyhemoglobin. *Proc Natl Acad Sci U S A* 22:210-6

Perlbarg V, Bellec P, Anton JL, Pelegrini-Issac M, Doyon J, Benali H. 2007. CORSICA: correction of structured noise in fMRI by automatic identification of ICA components. *Magn Reson Imaging* 25:35-46

Plummer C, Harvey AS, Cook M. 2008. EEG source localization in focal epilepsy: where are we now? *Epilepsia* 49:201-18

Prasad PV, Storey P. 2008. Magnetic resonance imaging. In *Molecular biomethods handbook*, ed. JM Walker, R Rapley, pp. 949-73. Totowa, NJ: Humana Press

Purves WK. 2004. *Life : the science of biology*. Sunderland, MA: Sinauer. 1 v. (various paging) pp.

Robson MD, Dorosz JL, Gore JC. 1998. Measurements of the temporal fMRI response of the human auditory cortex to trains of tones. *Neuroimage* 7:185-98

Semah F, Picot MC, Adam C, Broglin D, Arzimanoglou A, et al. 1998. Is the underlying cause of epilepsy a major prognostic factor for recurrence? *Neurology* 51:1256-62

Tao JX, Baldwin M, Hawes-Ebersole S, Ebersole JS. 2007. Cortical substrates of scalp EEG epileptiform discharges. *J Clin Neurophysiol* 24:96-100

Wiebe S, Blume WT, Girvin JP, Eliasziw M. 2001. A randomized, controlled trial of surgery for temporal-lobe epilepsy. *N Engl J Med* 345:311-8

Wilson FN, Bayley RH. 1950. The electric field of an eccentric dipole in a homogeneous spherical conducting medium. *Circulation* 1:84-92

Windischberger C, Langenberger H, Sycha T, Tschernko EM, Fuchsjager-Mayerl G, et al. 2002. On the origin of respiratory artifacts in BOLD-EPI of the human brain. *Magn Reson Imaging* 20:575-82

Worsley KJ, Liao CH, Aston J, Petre V, Duncan GH, et al. 2002. A general statistical analysis for fMRI data. *Neuroimage* 15:1-15

Worsley KJ, Marrett S, Neelin P, Vandal AC, Friston KJ, Evans AC. 1996. A unified statistical approach for determining significant signals in images of cerebral activation. *Hum Brain Mapp* 4:58-73

Yee SH, Gao JH. 2002. Improved detection of time windows of brain responses in fMRI using modified temporal clustering analysis. *Magn Reson Imaging* 20:17-26

Zhao X, Glahn D, Tan LH, Li N, Xiong J, Gao JH. 2004. Comparison of TCA and ICA techniques in fMRI data processing. *J Magn Reson Imaging* 19:397-402