Chapter 1. Introduction

In this chapter I introduce what we know about the function and anatomy of the brain areas discussed in this thesis (medial temporal lobe and cingulate cortex). I further discuss some features of epilepsy, with an emphasis on temporal lobe epilepsy and its treatment. All data presented in this thesis has been acquired from epilepsy patients. This thesis is not about epilepsy as such, but it is beneficial for the reader to understand the basics of epilepsy to better appreciate the clinical situation in which this research was conducted. The aim of this chapter is to set the stage for the results reported in this thesis so as to enable the reader to place them into context.

1.1 Memory

The capacity to learn and remember a seemingly infinite amount of information is one of the key facilities that makes us human. To illustrate this, ask yourself the following question: "Where where you on the following day: September 11th, 2001?". Chances are, you not only know where you were, but also how you heard about what happened that day, who told you about it, what you felt and what you thought was going to happen. This example illustrates the many components of memory — not only does this day mean something to you, but you can also retrieve a large amount of associated attributes. These attributes are not neutral facts. Rather, many of them have an emotional component. Thinking about the past does (introspectively) not just bring up a list of facts but rather each attribute is remembered with many of its significant autonomic attributes still attached. Have you ever met somebody, knew with high confidence that you knew the person but could not remember who the person is nor where you last met? This example illustrates a further example of memories: they are not all-or-nothing monolithic entities.

Rather, it is possible to retrieve some aspects of a memory (e.g., that it exists, “I have met this person before”) without any other of the attributes associated with it. Nevertheless the not-remembered information is often not lost — after some time, it might very well possible to retrieve the missing information.

What is most remarkable about memories is that they can be acquired very quickly. Most events for which we have memories happen only once and often last a very short time. This is nevertheless sufficient to build a representation that can last a lifetime. Life events are special in that they occur at a certain time to you personally (egocentric, i.e., they are episodic memories). This is distinct from other types of memories, such as memories for facts (semantic memories). Not only are fact memories not acquired from a single learning experience (usually one rehearses or studies facts) but also usually no attributes are associated with them — like, where and when did you learn this fact? Together, episodic and semantic memories build the class of all explicit long-term memories referred to as declarative memories. Another distinguishing feature of a declarative memory is a strong sense of confidence about whether one remembers something or not — performance and confidence are highly correlated (Bayley and Squire, 2002).

The other major class of memories are procedural, non-declarative memories. These include motor skills such as riding a bike which we can do effortlessly, but without the ability to articulate how we do it. That is, procedural memories are expressed by performance rather then recollection. Such memories require hundreds of learning trials (with feedback) to acquire. Procedural memories rely, for the most part, on brain structures distinct from those involved in declarative memories. These structures, such as the cerebellum and the basal ganglia, are not discussed here.

1.2 Anatomy, connectivity, and function of the medial temporal lobe

The neuronal structures necessary for the acquisition of declarative memories are situated in the medial temporal lobe (MTL). The MTL consists of a cortical and subcortical part. The cortical parts include the perirhinal, entorhinal, and parahippocampal cortices. Subcortically, the MTL includes the hippocampal formation (CA fields, dentate gyrus, subiculum) as well as the Amygdala (Squire et al., 2004; Squire and Zola-Morgan, 1991; Suzuki and Amaral, 2004). The MTL exhibits remarkable evolutionary consistency across several major mammalian species such as rodents, monkeys, and humans. While the absolute size differs, the gross anatomical and neurophysiological features are remarkably similar.

Destruction, inactivation or surgical removal of the MTL due to accidents, surgery, or stroke results in severe anterograde amnesia, manifested by profound forgetfulness. Some retrograde amnesia occurs as well, but memories of events that happened more than a few months before the injury are typically well preserved (but see below). Other cognitive capabilities are not impaired. In particular, short-term memory (working memory) is not impaired. Also, intelligence is not impaired. This is well illustrated by the well-studied patient H.M., who had bilateral MTL removal for treatment of epilepsy (Corkin, 2002; Milner et al., 1968; Scoville and Milner, 1957). H.M. has profound anterograde amnesia, but can remember events of his childhood. Also, he can learn new motor skills (procedural memory) such as mirror drawing but does not remember having done so. Even if the MTL is only deactivated temporarily, such as in global amnesia, the events that happened during this period are not remembered later.

The role of the MTL (and in particular the hippocampus) in memory is time limited. The loss of parts of the MTL causes temporally graded retrograde amnesia. In humans, such loss of

memory can be cover up to 15 years back in time (Corkin, 2002; Squire and Alvarez, 1995) . The extent of retrograde amnesia depends on how much of the MTL is damaged. For example, selective damage to CA1, an area of the hippocampus, resulted in 1–2 years of retrograde amnesia whereas more extensive damage can erase up to 15 years of memories (Rempel-Clower et al., 1996). Very remote autobiographical and factual memories beyond this time period appear unimpaired (Kirwan et al., 2008; Squire and Bayley, 2007). In animals, hippocampal damage induces retrograde amnesia lasting days to weeks only (rather then years in humans, see (Frankland and Bontempi, 2005) for a review). Thus, the hippocampus is not necessary for the retrieval of such memories. Rather, it is responsible for acquiring the memory. Over time, memories become independent of the hippocampus. While these facts are well established, it is not clear what the mechanisms are that make the hippocampus only necessary initially and lead to gradual independence. One framework originally formulated by Marr proposes that memories are first stored in the hippocampus and are then gradually transferred to cortical areas (Marr, 1970, 1971). He proposed that such transfer would occur through replay of activity during offline states (such as inactivity or sleep). While there have been several reports of such "replay" of activity (Buzsaki, 1998; Diba and Buzsaki, 2007; Foster and Wilson, 2006; Wilson and McNaughton, 1994), it remains to be demonstrated whether this indeed serves the purpose of transferring memories from the hippocampus to the cortex.

1.2.1 Anatomy of the hippocampal formation

The hippocampal formation consists of the hippocampus proper (cornu ammonis fields (CA), divided into CA3, CA2, and CA1) as well as the dentate gyrus (DG), subiculum,

presubiculum, parasubiculum, and entorhinal cortex (Andersen et al., 2007; Duvernoy, 2005). Note that in the human literature, the entorhinal cortex is often referred to as the "parahippocampal area". The hippocampal formation (hippocampus and dentate gyrus) is about 100x bigger in humans than in rats (rat 32mm^3, monkey 340mm^3, human 3300mm^3). Nevertheless, the basic anatomical features are remarkably similar between these 3 species. Not all areas show a similar increase in size from rats to monkeys (and humans). One noteworthy difference is the thickness of the pyramidal layer in CA1: it is about 5 cells thick in rats compared to 10–15 cells in monkeys. In humans, it is as much as 30 cells thick. It is estimated that the human hippocampal formation contains about 60 million neurons, compared to ~ 4 million in the rat (Andersen et al., 2007).

The hippocampus is tightly interconnected with the rest of the brain. It is distinct from cortical areas in that the connections with other brain areas are largely unidirectional. Cortical areas, on the other hand, are connected reciprocally (Felleman and Van Essen, 1991). Most input to the hippocampus first reaches the entorhinal cortex. From the entorhinal cortex, two major input pathways project to the hippocampus: the perforant path (to the dentate gyrus) and the temporoammonic alvear pathway (to CA1). The dentate gyrus does not project back to the entorhinal cortex (unidirectional). The dentate gyrus projects exclusively to CA3 (via the mossy fibers). All granule cells in the DG project to CA3. Their axons terminate in a stereotypical region of the CA3 cell body layer (stratum lucidum). CA3 pyramidal neurons send axons onto themselves (recurrent) as well as to CA1 (Schaffer collaterals). This circuit makes up what is referred to as the trisynaptic circuit. Synapse 1 is EC->DG, synapse 2 DG ->CA3, and synapse 3

is CA3->CA1. CA1 projects back to the entorhinal cortex as well as to the subiculum. There are no known direct connections between CA3/CA1 and the neocortex.

The entorhinal cortex receives input from a large number of cortical areas. The primate EC receives substantially more diverse cortical input than the rat EC. Prominent connections are from (and to) various high-level unimodal visual areas such as areas TE and TEO (through perirhinal cortex), area V4 (through parahippocampal cortex), numerous polysensory regions in the superior temporal gyrus, frontal areas such as the orbitofronal cortex, and the cingulate, as well as the insula. The EC also receives input from subcortical structures, such as the amygdala or the claustrum (see below).

1.2.2 Computational principles of hippocampal function

Due to the recurrent nature of CA3 pyramidal cell connectivity it has long been proposed that CA3 is the site of implementation of a large randomly connected recurrent network (Hasselmo et al., 1995; Kanerva, 1988; Marr, 1970, 1971; Rolls, 2007; Treves and Rolls, 1994). Such networks can be used to rapidly establish new attractors as well as for pattern completion (as demonstrated by Hopfield networks (Hopfield, 1982)). Pattern completion is crucial to implement content-addressable memories. However, the role of CA3 as a pattern completion engine has remained largely theoretical. A recent study using a genetically modified mouse strain that allows a selective and specific knockout of area CA3 function after animals reach adulthood reveals selective deficits that support this hypothesis (Nakazawa et al., 2002). Rather then being unable to learn at all, this mouse was unimpaired at a number of tasks such as the Morris water maze. There were no behavioral impairments in learning and retrieval of spatial memory.

Crucially, however, the mouse was unable to learn from single trials (such as a novel location of the platform) nor was it able to retrieve the platform location if only a partial set of the previous sets of cues were presented (Nakazawa et al., 2002; Nakazawa et al., 2003).

The dentate gyrus (DG), on the other hand, has been proposed to implement a complementary function: pattern separation (O'Reilly and McClelland, 1994; Rolls, 1996; Shapiro and Olton, 1994; Treves and Rolls, 1994). This suggestion is motivated by the following observations: i) Connectivity between DG dentate cells and CA3 pyramidal cells is very sparse. This results in a small degree of divergence (~ 14 pyramidal cells per granule cell) (Acsady et al., 1998). ii) There are large differences in the number of neurons in the EC (200'000), DG (1'000'000) and CA3 (300'000). All estimates are for the rat (Amaral et al., 1990; Amaral and Lavenex, 2007; Boss et al., 1985; Henze et al., 2002) but similar proportions are valid for the monkey and human DG (Amaral and Lavenex, 2007; West and Slomianka, 1998). This leads to an expansion (EC->DG) followed by a contraction (DG->CA3) of effective dimensionality (see below). Thus, every CA3 neuron receives (with high probability) input from a different subset of DG granulate cells. iii) Only a small fraction of granule cells is activated in any given task (Chawla et al., 2005; Jung and McNaughton, 1993; Witter, 1993). From theoretical studies it is known that such a projection effectively increases the distance between every possible pattern, thus making the patterns more dissimilar to a downstream region. This is due to the increase in the effective dimensionality. The powerful computational properties of such a construct are well demonstrated by "liquid state machines", which essentially consist of random sparse projections from a low-dimensional space into a high-dimensional space and back (Maass and Markram, 2004; Maass et al., 2002). Several experimental studies support this hypothesis. Until recently,

the only experimental support for this hypothesis has come from behavioral observations after DG lesions in rats (Gilbert et al., 2001). In match-to-sample tasks, DG lesioned rats show deficits in distinguishing different objects (some of which indicate food and others don't) if they are close together in space. They have little deficits if objects are far apart (spatially). A more direct demonstration of the role of the DG in pattern separation comes from genetic NR1 knockouts restricted to DG granule cells (McHugh et al., 2007). Since NR1 is a necessary subunit of the NMDA receptors, this mutation prevents NMDA dependent plasticity in granule cells. No perforant path (EC->DG) potentiation could be invoked in these mice. Behaviorally, these mice had difficulty distinguishing between different contexts as measured by inappropriate freezing in contexts which are similar to, but slightly different, from the context in which conditioning took place. Mice with intact DG had no problem in distinguishing the two contexts. Interestingly, the deficit was only temporarily: more training (experience) could overcome the deficit. Thus, pattern separation is important for the rapid acquisition of new experiences. In spatial tasks, DG granule cells have place fields which are similar to CA3 and CA1 pyramidal cells. Place fields rapidly change their firing preference if the external environment (size, color, shape) is changed ("remapping"). Interestingly, remapping of DG place fields occurs more rapidly (for smaller environmental differences) than for CA3 cells. This is demonstrated by the finding that correlations between the firing activity of populations of DG cells decay rapidly as a function of small changes of the environment, whereas CA3 cells decorrelate only after large changes (Leutgeb et al., 2007).

These experimental findings (for both CA3 and DG) are the first to offer direct support for the long-standing theoretical proposal that one of the functions of the hippocampus is pattern separation (DG) followed by pattern completion (CA3).

1.2.3 Amygdala

The amygdala receives direct inputs from all sensory systems as well as the hippocampus (Aggleton, 2000). Its projections, amongst others, to different areas in the hypothalamus and brain steam and can thus directly influence the autonomic nervous system. The amygdala consists of several separate nuclei (central, basal, lateral, and medial). Each nucleus is either mainly an input or an output structure. The lateral nucleus receives input from all sensory systems. The central nucleus projects to the brainstem and hypothalamus, whereas the basal nucleus projects to cortical areas as well as the striatum. There is very substantial reciprocal connectivity between the hippocampus and the amygdala. The lateral and basal nuclei of the amygdala project prominently to the entorhinal cortex. Feedback connections from the EC terminate mostly in the basal nucleus of the amgydala. The amygdala is necessary for some kinds of rapid learning such as Pavlovian fear conditioning (Fanselow and LeDoux, 1999). Forms of synaptic plasticity like long-term potentiation (LTP) can be induced both in vivo and in vitro (Chapman et al., 1990; Rogan et al., 1997). While the amygdala is not necessary for declarative memory formation, it nevertheless modulates the strength of such memories (Phelps, 2004; Richardson et al., 2004). The amygdala is thus crucially involved in the acquisition of some types of memories. In the following chapters I will show that many of the single-neuron responses related to single-trial learning can be observed similarly in both the amygdala and the hippocampus.

1.2.4 Adult neurogenesis

Most neurons in the adult brain are postmitotic. Remarkably, there are a few exceptions: there are progenitor cells (stem cells) in the subgranular and subventricular zone that continuously divide and send new neurons to the adult dentate gyrus and the olfactory bulb, respectively. There are indications that other regions of the brain contain new neurons as well (Garcia et al., 2004; Gould et al., 1999). This adult form of neurogenesis occurs throughout life and has been shown to be modulated by numerous environmental factors such as stress, learning and exercise (van Praag et al., 1999; van Praag et al., 2002). Of interest to the results of this thesis, some report a relationship between the sensitivity to novelty and the number of new neurons in the dentate gyrus (Lemaire et al., 1999). The discovery of postnatal neurogenesis in the MTL suggests the intriguing possibility that it is related to our capacity to learn. It will be very interesting to explore the computational implications of the incorporation of new neurons into existing circuits, because new neurons have very distinctly different electrophysiological properties (van Praag et al., 2002), such as increased plasticity (Schmidt-Hieber et al., 2004). It is unclear how these different single-cell properties affect circuit function (Lledo et al., 2006). Evidence has recently accumulated that inappropriate incorporation of new neurons is implicated in epiloptogenesis (Buhl et al., 1996; Covolan et al., 2000; Parent et al., 1997). Whether this is a cause or an effect, however, is unclear, but it is a very promising route for further experimental investigation (Parent, 2007).

1.3 Mechanisms of plasticity—Circuit and single-cell properties

Establishing new memories (learning) is thought to require structural changes (plasticity) for long-term storage (Martin et al., 2000). Here, any mechanism that changes the composition, shape, size, or configuration of a cell is referred to as structural plasticity. Examples are insertion or removal of proteins such as neurotransmitter-activated ion channels, the removal or growth of new spines, as well as the growth of new axons. This form of memory is distinct from activity-based memories such as working memory, iconic memory, adaptation, or priming which (at least in theory) do not require permanent structural changes. Such memories only last as long as the neuronal activity (in the form of spiking or subthreshold processes such as deactivation) remains active—typically only a few seconds (see (Wang, 2001) for a review and (Romo et al., 1999) for an example of pre-frontal working memory activity). Longer-lasting memories, however, do not require constant activity. We do not lose our memories after we sleep, go into deep anesthesia, or have a severe epileptic seizure. A mechanism must exist that transforms information about the external environment (represented as neural activity) into changes in the brain. The exact nature of these changes is a matter of great debate and is largely unknown (see below). Understanding this mechanism at all levels involved has been one of the goals of neuroscience research since its beginning (Bliss and Lomo, 1973; Hebb, 1949; Pavlov, 1927). The results presented in this thesis contribute to this understanding by demonstrating that there are single neurons in the brain that function as generic novelty/familiarity detectors. It is hypothesized that these detectors are part of the system that initiates learning for a novel stimulus.

Long-term potentiation (LTP) and long-term depression (LTD) are a class of molecular and cellular mechanisms that can trigger such long-lasting changes in synaptic strength (Bliss and

Lomo, 1973). Here, I will briefly summarize what we know about LTP as well as its relevance to behavioral changes related to learning.

The amount of postsynaptic current influx evoked by presynaptic release depends on many different factors, such as the number of vesicles released, the number of channels in the postsynaptic terminal, internal Ca2+ stores, and the number, types, and composition of voltage-dependent channels (among many others). The modification of any of these factors potentially leads to changes in synaptic strength, i.e., the degree and amount of influence presynaptic release has on the postsynaptic neuron. Similarly, the number of synaptic contacts between two neurons can change as well (formation and destruction of synapses). It has been observed consistently that the synapses connecting two neurons get strengthened when the presynaptic neuron fires shortly before the postsynaptic neuron (LTP). Reversal of the temporal order (postsynaptic neuron fires before presynaptic) leads to a decrease of the synaptic weight (LTD). The existence of LTP/LTD has by now been shown in a wide variety of species and brain structures including the hippocampus, amygdala, and the neocortex. This is one of the fundamental principles of synaptic plasticity— loosely summarized "what fires together wires together". Hebb originally postulated that "when an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased" (Hebb, 1949). LTP/LTD is a candidate mechanism that implements this rule (see below).

This principle (Hebb's law or Hebbian learning) is commonly expressed as a correlation-based learning rule that describes the incremental change of the synaptic weight as the product of the pre and postsynaptic firing rates: $\Delta w = v_j v_i$. Since this implies that weights can grow

infinitely, a maximum weight $w = \min(w + \Delta w, w_{\max})$ is usually imposed. Also, in this simple form, weights never decrease (see below for more detailed discussion). This principle can be observed at many different levels of organization, starting with a single synapse between two neurons all the way to behavioral observations (for example, Pavlovian conditioning). While we have some understanding of the detailed molecular mechanisms, the intermediate levels are much less clear and are poorly understood. For example, during Pavlovian conditioning, how is it that plasticity can be induced selectively at exactly the right synapses while not influencing all the other existing synapses? After all, acquisition of a new memory does not require that an existing memory be overwritten. Also, whether the mechanism of LTP is sufficient and/or necessary for learning of new memories remains to be demonstrated (Martin et al., 2000; Shors and Matzel, 1997; Stevens, 1998). One strategy to demonstrate that LTP is indeed sufficient for learning would be to monitor a large number of synapses during memory acquisition. If reverting these very same synapses after learning back to their original strength erases the memory it is demonstrated that these synaptic changes did indeed result in the observed behavioral change (Neves et al., 2008). A large number of studies have been conducted that attempt to demonstrate that LTP is indeed the mechanism underlying learning (see (Martin et al., 2000) for a review). For example, it has recently been demonstrated that behavioral single-trial learning of inhibitory avoidance by rats induces molecular and electrophysiological changes in the hippocampus very similar to those induced by artificial LTP induction in CA1 (Whitlock et al., 2006). While these and others suggest a close link between LTP and learning, this has not been convincingly demonstrated at this point of time.

The amount and direction of plasticity at a single synapse depends on many factors (see above). One of the most important factors, however, is time. For a long time, it was not clear what exactly qualifies as “fires together”. What qualifies as “together” ? Within 1 sec or within 1 ms ? Also, does the order matter (first presynaptic, then postsynaptic or the opposite)? One of the fundamental principles that has emerged is that both the temporal order as well as temporal distance matter on the order of milliseconds. Technical advances allowed the first experimental demonstration by dual intracellular patch recordings from two synaptically connected neurons (Bi and Poo, 1998; Markram et al., 1997). The remarkable finding was that evoking a spike in the presynaptic neuron 10 ms before evoking a spike in the postsynaptic neuron lead to strengthening of the synapse. The reverse (presynaptic spike follows the postsynaptic spike) leads to weakening of the synapse. Spacing the spikes closer together in time evokes stronger changes. Spikes that occur too far apart in time (> 40 ms) fail to induce any changes in synaptic strength. This mechanism is referred to as spike-timing-dependent plasticity (STDP). The time window for induction of STDP is at best ± 40 ms. In classical conditioning, the conditioned stimulus (CS) can be separated by up to several seconds from the unconditioned stimulus (US) (Pavlov, 1927). This is thus much bigger then the timescale over which STDP can occur. Additional mechanisms (such as working memory) must thus exist to bridge this time gap. Many computational learning algorithms such as reinforcement can be implemented with STDP as a mechanism. Such learning is usually referred to as correlation-based learning (or “Hebbian learning”). The existence of STDP has been shown in a large number of different species, brain areas, and cell types (see (Caporale and Dan, 2008) for a review).

The induction of long-lasting plastic changes is often dependent on changes in intracellular calcium (Ca^{2+}). There are many different types of neurotransmitter or voltage-gated ion channels which are permeable to calcium. This accumulation of intracellular Ca^{2+} triggers many molecular events which eventually lead to long-lasting changes of synaptic strength. Without calcium influx, LTP can not occur. This is convincingly demonstrated by the absence of LTP if the NMDA channels are blocked pharmacologically during LTP induction. The induction of plasticity itself is not sufficient to ensure a long-lasting structural change. If certain molecular processes are disrupted, the change does not last long. For long-lasting LTP ("late LTP"), gene transcription and the synthesis of new proteins is required (Huang et al., 1996; Kelleher et al., 2004; Schuman, 1999; Squire, 1992; Sutton and Schuman, 2006). This is also true in vivo: memories do not last if synthesis is inhibited (Davis and Squire, 1984; Flexner et al., 1963; Squire, 1992). The blockage of protein synthesis (during induction) thus prevents the conversion of a short-term to a long-term synaptic change. Rather, it decays back to baseline within a few hours. Synaptic changes can, however, occur without new proteins, i.e., the modification and relocation of existing proteins is sufficient. While protein synthesis blockers present during the induction of LTP only affect the late phase of LTP using traditional LTP induction protocols, there are other induction protocols for LTP where protein synthesis is required also for the early phase. Application of neurotrophic factors such as BDNF can induce potentiation without electrical stimulation (Kang and Schuman, 1995; Levine et al., 1995; Lohof et al., 1993). Application of synthesis inhibitors prevents this kind of induction of LTP (Kang and Schuman, 1996). There are also situations where synthesis is required for early LTP, such as high synaptic

background activity (Fonseca et al., 2006). In the hippocampus, synthesis can also be modulated by changes in extracellular dopamine (Smith et al., 2005).

1.4 Temporal lobe epilepsy

Epilepsy is one of the most common forms of neurological impairment. The lifetime risk of experiencing at least one seizure is 3% (Chang and Lowenstein, 2003). About 1% of all people develop unprovoked seizures without any obvious reason (Steinlein, 2004). In the USA, an estimated 1.1 - 2.3 million people have epilepsy. Epilepsy is a medical condition characterized by the presence of recurrent seizures. Clinical manifestations of seizures can include loss of consciousness, involuntary twitching of muscles, brief periods of amnesia, sleep disturbances as well as other sensory, cognitive, psychic, or autonomic disturbances.

Seizures are fundamentally a circuit-level phenomenon. Thus, the study of seizures requires a systems perspective. They are thought to occur due to hypersynchronous neuronal discharges that lead to uncontrolled spread of excitatory activity to other areas of the brain. Any complex neuronal circuit relies on a tight balance between inhibition and excitation to function properly. This is particularly true for the cortex as well as the hippocampus due to extensive recurrent excitation. Reasons for synchronous discharges can either be loss of inhibition, an increase of excitation, or a mixture thereof. In some cases, the causes of epilepsy are clearly attributable to a specific component of the circuit, such as specific genetic mutations of voltage-gated ion channels. Such mutations have been identified for potassium (K), sodium (Na+) and chloride (Cl) channels. These mutations directly affect the excitability of the circuit. One

particular example is a loss-of-function mutation in the Na+ channel Beta1 subunit (SCN1B) which leads to slower inactivation and thus more current influx (Wallace et al., 1998).

Epileptic seizures are classified on several different dimensions. Complex seizures result in loss of consciousness whereas simple seizures do not. Generalized seizures arise simultaneously in the entire brain (bilateral) whereas partial seizures arise from a local (unilateral) area of the brain. Partial seizures can spread and progress to a generalized seizure (secondary generalization). Generalized seizures include absence (petit mal) and grand-mal (tonic-clonic) seizures. Absence seizures are special because they are very brief and occur without warning. Patients cease normal activity and stare for only a few seconds and then return to normal immediately afterwards. They have no memory of the epileptic episode. These events can occur hundreds of times a day.

One of the most common forms of epilepsy in humans is temporal lobe epilepsy (TLE) (Engel, 2001; Ojemann, 1997). TLE seizures are usually complex partial and thus manifest themselves with alteration of consciousness. Often, a simple partial seizure precedes the complex partial. This phenomenon is referred to as an aura, as patients get a physical awareness of the onset of seizure before it progresses to cause an alteration in consciousness. In contrast to other types of epilepsies, medial temporal lobe epilepsies are often difficult to control with antiepileptic drugs (AEDs). It has been speculated that this might be due to biased criteria for the pre-clinical evaluation of drug candidates, which are usually screened only for effectiveness for petit mal absences and tonic-clonic seizures. As a result, about 50% of patients (Ojemann, 1997) require other treatment to achieve control of their seizures. For many, surgical removal of the epileptogenic parts of the MTL is an option. 70–90% of patients become free of disabling

seizures after surgical treatment (Engel, 2001). This indicates that an underlying cause of TLE is a structural abnormality restricted to the MTL. Post-surgical histology of the removed tissue often shows marked loss of principal neurons of the hippocampus (sclerosis).

Planning for surgery requires extensive preparation to evaluate the location and extent of the resection, as well as an evaluation of possible loss of function resulting from the resection. Accurate localization can be very difficult and time consuming. This is particularly true for one of the most common pathologies that results in temporal lobe epilepsy: hippocampal sclerosis. The neuronal loss in the hippocampus is hard to detect using conventional structural MRI (see below) unless it is severe.

Several non-invasive indicators can be used to identify possible seizure origin areas: structural magnetic reasonance imaging (MRI), computed tomography (CT), positron emission tomography (PET), surface electroencephalography (EEG), or the behavioral symptoms accompanying a seizure. Only about 30% of all seizures are caused by tumors or lesions that are visible on MRI or CT (Engel, 2001). Others can be identified from surface EEG recordings of multiple seizures. If these methods fail to clearly localize the seizure (or if the methods contradict each other), invasive recording techniques can be used (Spencer et al., 2007). These include subdural grids of electrodes placed on the surface of the cortex as well as depth electrodes (see below). Together with video monitoring and surface EEG, such recordings allow accurate and high-resolution tracking of the evolution (pre-ictal, ictal, post-ictal) of a seizure (the "ictal" event) Intraoperative recordings on the surface of the cortex typically only allow the recording of interictal spikes (epileptoform EEG) but not spontaneously occurring seizures. It is thus frequently necessary to implant semi-chronic electrodes to record activity continuously until a

seizure occurs. Electrodes are implanted bilaterally at likely epileptogenic sites using a lateral approach. We always implanted electrodes in the hippocampus, amygdala, anterior cingulate cortex, orbitofrontal cortex, and supplementary motor cortex. The exact implantation site was determined based on clinical criteria with the help of co-registered CT, structural MRI, and angiogram. Implantation was guided by a stereotactic frame fixed to the head of the patient (Spencer et al., 2007). Monitoring can take up to several weeks of continuous observation and recording. If activity preceding a seizure has a clear unilateral origin at a restricted set of electrodes (for example, the anterior hippocampus) surgical removal of that part of the brain controls seizures in a large fraction of cases (70–90%) with little or no functional deficits (Spencer et al., 2007; Vives et al., 2007). Since activity has to be recorded continuously (while waiting for a seizure to occur), there are large periods of time where brain function is normal but intracranial signals can be recorded. This gives us (scientists) the unique opportunity to directly observe the electrical activity of the awake human brain during behavior. The data reported in this thesis are all recorded during these periods of time.

It should also be mentioned that as non-invasive diagnostic technologies improve, the need for invasive electrode implantation will decrease. Thus, the window of opportunity to record from epilepsy patients for research purposes could eventually close. While this seems far into the future one should nevertheless keep this in mind when planning a new research program on the approach described here. There are, however, several other surgeries such as deep brain stimulation (DBS) implantation and small resection for treatment of severe psychological problems such as OCD (Williams et al., 2004) that will open up new opportunities to apply this approach.

1.5 Electrophysiology in epilepsy patients

1.5.1 Clinical

For clinical electrophysiology, two primary types of electrodes are used: subdural grids/strips and depth electrodes. Depth electrodes have 4–12 regularly spaced contacts (2–6 mm spacing) along the entire length and can thus be used to record intracranial EEG along the entire depth of the cortex and subcortical structures if implanted perpendicular to the cortex. Grid electrodes have arrays of disk electrodes (3–4 mm diameter) imbedded in a thin sheet of plastic so that the uninsulated side of the electrode rests on the pial surface of the cortex (Wilson, 2004).

1.5.2 Research

Two types of signals acquired from epilepsy patients with implanted electrodes can be used for research purposes. First, the signals originating from the clinical contacts (on grids as well as depth electrodes) can be utilized to record low-frequency (typically < 100Hz) local field potential (LFP). Due to their low impedance (< 1 kOhm) and large size, however, these electrodes do not allow the recording of local, small and fast extracellular currents such as those evoked by spikes. Second, microwires embedded in the depth electrode (the so-called “hybrid depth electrode”, AD-Tech Medical Instrument Corp, Racine WI), can be used to record well localized LFP as well as single-unit activity. The wires are made of isolated Platinum/Iridium and are 40 μm in diameter (Fried et al., 1999). We used electrodes with 8 embedded micro wires. Their tip is exposed by cutting the wires to appropriate length (5 mm typical) during surgery. Due to their much smaller exposed tip, microwires have much higher impedance (several 100–500

kOhm at 1 kHz; see methods for detailed measured values). This allows the measurement of single-unit activity with high reliability.

1.5.3 Previous human single-neuron studies

Recording from single neurons in awake, behaving humans offers the unique opportunity to address questions about the function and structure of our brains– addressing questions that have proven difficult or impossible to address using animal models. While it is possible to investigate a large number of mechanisms using animals, there are capabilities which are either very difficult to assess in animals or are unique to humans. These include language, episodic memory, rapid learning, emotions, remote memory, planning, and subjective experience. In the following I will review what has been discovered so far using single-cell recordings from humans (Engel et al., 2005; Kreiman, 2007; Wilson, 2004), with an emphasis on studies reporting findings that would have been hard or impossible to achieve with animals. Also, the focus will be mainly on learning and memory and thus on medial temporal lobe recordings. Such recordings are predominantly from epilepsy patients who are being evaluated for surgery. The many studies reporting single-cell recordings from sub-corticial structures such as the subthalamic nucleus or the thalamus will not be reviewed here. These recordings are made interoperatively from Parkinson's patients and are mostly related to motor responses rather then memory.

Human in-vivo single-unit recordings have a long history. The earliest recordings were made interoperatively before resection of tissue (Verzeano et al., 1971; Ward and Thomas, 1955). Some of the earliest comprehensive studies that used semichronically implanted depth electrodes in the MTL already identified neurons that are selective to specific words, faces, stimulus

on/offset, or motor responses (keypresses) (Halgren et al., 1978a; Heit et al., 1988, 1990). This has started a long debate as to whether these responses are visual or memory responses. The ventral visual stream contains consecutively more invariant and selective neurons that respond to very abstract concepts with the highest level of abstraction observed in the inferior temporal cortex. It is thus conceivable that responses in the hippocampus and closely connected cortical areas represent a continuation of such responses. On the other hand, they could be completely different in that high-level ventral stream responses are the fixed "vocabulary" of very well-known entities (such as animals, cars, trees) and MTL responses reflect whether these particular objects had been seen before or not. This strict dichotomy between memory and visual responses seems somewhat artificial, however. After all, any visually selective response that is not genetically innate is a "memory". I propose that a more natural way to look at this distinction is as a continuous gradient of recency: abstract, long-term, and stable category-type memories (like "animal") are represented in the ventral stream (inferior temporal areas such as TE and IT in monkeys) whereas more recent or specific memories such as "this particular animal" are represented in the MTL. This gradient could continue even further down the ventral stream to areas such as V4, where very long-term knowledge about basic visual features (such as colors) is represented (Gallant et al., 1996; Gallant et al., 2000). It is conceivable that such responses are plastic as well over the long-term. As a memory becomes more permanent and abstract (such as learning a new category), responses gradually emerge in the inferior temporal lobes. Available data about object selectivity and its emergence in non-human primates is well compatible with this view.

Studies by Fried et al. revealed a much better understanding of the responses evoked by presentation of visual stimuli such as objects and faces (Fried et al., 2002; Fried et al., 1997). First, neurons were identified that distinguished between faces and objects. However, neurons also responded selectively to attributes of faces (such as gender and the emotions happy, surprise, fear, disgust, angry, sad, and neutral). Additionally, many neurons respond differently when being exposed to a stimulus that has never been seen before by the patient (novel) compared to a familiar stimulus. Further recordings by Kreiman et al. revealed that a widespread feature of neurons in the amygdala, hippocampus, and entorhinal cortex is the selectivity to visual categories (Kreiman et al., 2000a). Such neurons show a highly invariant response to any instance of a broadly defined visual class such as animals, cars, objects, or faces. Also, these neurons follow the actual percept rather then the physical (retinal) input in a rivalry design (Kreiman et al., 2002; Reddy et al., 2006). Even when patients are asked to imagine a previously seen picture (such as of the well-known politician Bill Clinton), the same neurons respond both to a picture of Bill Clinton and to an imagined image of Bill Clinton (Kreiman et al., 2000b). Thus the response of these neurons follows the actual visual percept rather then the physical input. Extending this finding, the group of Koch et al. identified neurons which are highly selective (such as for a particular person) as well as highly invariant (Kraskov et al., 2007; Quiroga et al., 2007; Quiroga et al., 2005; Waydo et al., 2006). That is, their response is sparse in the sense that any given neuron only responds to a very small subset of all tested stimuli (response sparseness). Other studies of single-unit activity in relation to learning, particularly comparisons between viewing a stimulus the first and second time, are discussed in the following chapters (Cameron et al., 2001; Viskontas et al., 2006).

Another line of investigation that has benefited greatly from human single-unit recordings has been the study of language (Creutzfeldt et al., 1989a, b; Ojemann et al., 1988; Ojemann et al., 2002). Taking advantage of the possibility to record from temporal cortex that is later resected, Ojemann et al. has recorded a large variety of responses evoked by language comprehension and production. While not part of the MTL, the lateral temporal cortex is known to be crucial for declarative memory of verbal material (Ojemann et al., 1988; Ojemann and Dodrill, 1985; Perrine et al., 1994). Language function is strongly lateralized in the dominant hemisphere and the ability to record from both the dominant and non-dominant hemisphere has contributed to this understanding. Interestingly, only few neurons sampled from the surface of the lateral temporal cortex respond to visual stimulation (in contrast to the MTL responses summarized above). Rather, neurons responded to either silent (reading without pronouncing) or overt speech. Other neurons responded to the memorization or retrieval of verbal memory material. Superior temporal gyrus neurons respond very prominently while subjects listened to spoken language and were often selective to specific combinations of consonants (Creutzfeldt et al., 1989a). Also, neurons were found that respond preferentially to the patient's own voice.

A series of electrical stimulation studies by Halgren et al. revealed crucial insights into the results of direct temporary disruption of the MTL (Halgren et al., 1978b; Halgren and Wilson, 1985; Halgren et al., 1985). While it was well known that bilateral structural damage caused severe amnesia, the causal functions of the MTL can only be established by selective (and reversible) disruption. This has been achieved by injecting a short pulse (100 μs) of current into a number of depth electrodes simultaneously with the onset of a novel or repeated stimulus (Halgren et al., 1985). This single pulse of stimulation disrupted normal ongoing activity for 400

ms. Patients where shown a series of stimuli and had to indicate whether the stimulus had been shown before or not (old/new). Stimulation was applied either during learning, retrieval, or both. This stimulation protocol did not disrupt performance if the delay between presentation of the two stimuli (new and old) was short (2 s). Performance was severely impaired if stimulation was applied during both learning and retrieval (16% correct vs. 66% without stimulation). Most interestingly, performance was severely impaired if stimulation was selectively applied either during learning (no acquisition) or retrieval. The same stimuli that could not be retrieved if stimulated during retrieval could be retrieved if there was no stimulation. This demonstrates a direct causal role of the human MTL in both memory acquisition as well as retrieval (at least of relatively recent memories). It also demonstrated, together with other studies (Chapman et al., 1967; Halgren and Wilson, 1985; Ojemann and Fedio, 1968), that MTL stimulation does not disrupt perception, decision making, response execution, retrieval of remote memories, or otherwise severely alter the cognitive state. It is also interesting to note that patients reported with confidence that they had not seen the stimulus before and not that they did not know or could not answer the question (Halgren and Wilson, 1985).

Electrical stimulation of the temporal lobe (using similar techniques as described above) during the absence of external visual input can evoke a series of phenomena such as deja-vu (a strong sense of familiarity), complex hallucinations, alimentary sensations, fear or anxiety, and amnesia (Bancaud et al., 1994; Halgren et al., 1978b; Penfield, 1958; Penfield and Perot, 1963). Authors have also remarked on the extreme variability of the type of effects evoked (stimulation sites and patients). This confirms that the temporal lobes are directly involved in the retrieval of memories.

1.6 The anterior cingulate cortex

The cingulate cortex is located directly above the corpus callosum (Allman et al., 2001; Paus, 2001). Its anterior part is referred to as the anterior cingulate cortex (ACC). It is thought to play a crucial role in many higher cognitive functions such as error monitoring, attentional control, conflict resolution and reward processing. Similarly, ACC dysfunction has been attributed to several major pathologies such as obsessive-compulsive disorder (OCD), bipolar affective disorder (BAD), chronic pain, and major depression. Surgical removal of parts of the ACC has proven to be a successful treatment of last resort for these pathologies (Jung et al., 2006a; Williams et al., 2004). The ACC also plays a major role in cue-induced craving in drug addiction (Kalivas and Volkow, 2005). Despite this, the function of the ACC remains poorly understood and controversial. One of the common elements of the above pathologies is a major learning deficit. In OCD, for example, inappropriate actions are repeated over and over despite explicit knowledge of their negative consequences. Similarly in cue-induced craving, drug consumption is induced because of a strong association between a cue and rewarding behavior.

A variety of functions have been attributed to the ACC; however, it has proven difficult to study with animal models (lesions, electrophysiology; but see (Frankland et al., 2004; Han et al., 2003)). While this is certainly partly due to the poorly understood function of the ACC, another likely reason is that the ACC is important for cognitive functions that are difficult or impossible to study and quantify in animals. Additional difficulty is added in that the non-human primate homologue of ACC is not clearly anatomically defined and overlaps with the cingulate motor area (CMA). In non-primate mammals the anterior cingulate exists but the anterior part (subgenual) is referred to as the prelimbic and infralimbic areas (Uylings et al., 2003). Also, the

ACC (Brodmann's Area 24) contains especially large spindle cells ("von economo neurons") that exist only in humans and great apes, but not other mammals including monkeys (Allman et al., 2001; Nimchinsky et al., 1999). This suggests that parts of the function of the ACC might be unique to humans and very closely related species.

Here I synthesize the conclusions from a number of studies and critically review what is known about the function of the ACC. I will start by reviewing what has been learned from event-related potential studies and will proceed by discussing how these findings have been extended and revised based on fMRI studies. Finally, I will compare those findings to human lesion studies and point out a number of discrepancies with the fMRI studies.

1.6.1 Reward and Dopamine

Dopamine (DA) is a crucial modulator of learning. It is released by DA neurons located in the ventral midbrain and the striatum. One of the prominent projections of midbrain DA neurons is the ACC (Gaspar et al., 1989). These DA neurons fire in short bursts in response to unexpected rewards, the expectation of reward, as well as novel items (Schultz, 2000). Similarly, many drugs of abuse induce the release of massive amounts of dopamine. Learning-induced plasticity such as long-term potentiation (LTP) is profoundly modulated by the presence of DA. Changing behavior in response to external feedback such as reduced reward is signaled by single neurons in the monkey ACC (Shima and Tanji, 1998). The reversal learning task requires subjects to associate arbitrary stimuli with actions. From time to time, the reward contingencies reverse. Thus, the optimal behavior in this task is to switch to the other action in response to receiving reduced reward. Similarly, if there is no error, the beneficial action should be sustained.

Deactivation of the ACC by lesions (Kennerley et al., 2006; Williams et al., 2004) or temporary deactivation (Shima and Tanji, 1998) profoundly impairs performance in this task in both monkeys and humans.

1.6.2 Event-Related Potentials

Before the advent of fMRI studies, the function of the ACC had mostly been studied with event related potentials (ERPs). Using reaction time (RT) tasks that require a fast response to sometimes conflicting stimuli, subjects make errors. Subjects are usually aware that the response was an error before the feedback signal. An example of such a task is the Stroop interference task (Kerns et al., 2004): subjects are shown a word on a screen, printed in a particular color. The task is to respond, as fast as possible, by pressing a button that indicates the color the words are shown in. For example, the word could either be red or green. These two words are printed either in red or green. Of the 4 possible combinations, 2 are congruent (red, green) and 2 incongruent (red, green). Incongruent trials require a significantly longer time to respond than do congruent trials (typically 40–80 ms on average). Additionally, if speed is more important than accuracy, incongruent trials evoke more erroneous responses.

A prominent observation during such tasks is a negative potential referred to as the error-related negativity (ERN). The potential peaks over frontal-parietal electrodes at 100 to 150 ms after the response (Paus, 2001). But because ERPs are recorded on the scalp, it is impossible to localize the source of the signal precisely. However, dipole models can be utilized to predict possible configurations of electric sinks and sources that could account for the data. One possible

model is a dipole located in the ACC. These observations motivated a number of imaging studies that attempted to definitively localize the source of this potential.

1.6.3 Neuroimaging studies: PET and fMRI

Early imaging studies, comparing blocks of trials between which only the variable of interest changed ("block design"), have suggested that the ACC is generally involved in the executive control of cognition. These studies have used tasks that require selective attention, working memory, and self-monitoring for errors (Bush et al., 1998). None of these tasks was found to elicit ACC activity that was specific to a particular function and it was thus proposed that the ACC generally responds to task difficulty, irrespective of the specific task. PET studies have generally reached the same conclusion (Paus et al., 1998): ACC activity is most strongly correlated with task difficulty. Since these results are achieved by subtracting the activity of two different blocks (easy and difficult), task difficulty refers to any possible variable that can make a task more difficult. This, for example, includes different demands for working memory, attention, difficulty of cognitive analysis (e.g., reading versus telling the font), and increased demand for motor commands (e.g., precision of movement). Because many experimental manipulations have been observed to change ACC neural activity, a number of different hypotheses for ACC function have been advanced. Two influential theories are the "conflict monitoring" and the "error detection theory". It has, however, proven difficult to clearly disambiguate the predictions that different theories make by using blocked designs.

1.6.4 Conflict monitoring, error monitoring, and cognitive control

One of the prevalent views is that one of the functions of the dACC (dorsal part of the ACC) is conflict monitoring. This interpretation was primarily developed because of earlier blocked studies of the Stroop interference task that showed enhanced activation of a part of the dorsal ACC when comparing blocks with and without interference effects. However, interference can cause multiple effects. Firstly, it creates a conflict between a fast, overtrained automatic process (reading the word) with a slower process (telling the color of the ink the word is printed in). Secondly, the very same effect increases attentional load, task difficulty, and behavioral errors. It is thus important to attempt to dissociate between these effects. One possibility is to use an event-related design where the degree of interference is modulated separately from task difficulty. One study (Carter et al., 1998) applies this approach. Carter et al. use a task where the subject is first presented with a cue (A or B) and than later with a probe (X or Y). The subject is instructed to only respond if the cue is an A and the probe an X. All other combinations are presented with low probability. They can be divided into low (BY) and high (BX,AY) interference trials. Task difficulty was modulated by removing pixels from the cue and probe stimuli (harder to read). The authors found that activity within ACC was higher for interfering vs. non-interfering trials (compatible with conflict monitoring hypothesis). However, they also found increased activity for error trials vs. correct trials. Interestingly, the authors found that ACC activity was also increased (relative to baseline) for the correct and non-interfering trials, but less than in error or interference trials. This finding resolves the previous debate in demonstrating that ACC is actually activated by both error and interference. However, this finding does not rule out that both types of activation are the result of a common underlying cause. There are multiple

possibilities that could cause this pattern of activity without any relationship to conflict or error monitoring. One example is increased attentional load, which could be caused both by an error as well as by interference. Also, errors are not independent from interference: more errors are made if interference is higher. The conclusions that can be drawn regarding the function of the ACC from these types of design are thus limited.

Two follow-up studies (Botvinick et al., 1999; Kerns et al., 2004) have shed light on this issue by using a Stroop interference task (color naming, see above). In an attempt to disambiguate conflict monitoring from error-related activity, the authors compared BOLD ACC activity of incongruent trials (I) which were either followed by a congruent (c) or an incongruent (i) trial (cI or iI). If ACC activity is related to interference itself, activity on both types of trials should not differ. If, however, ACC activity relates to the monitoring of conflict and the subsequent induction of control, activity on iI trials should be lower than activity on cI trials. Also, the reaction time for the I trials which follow an incongruent trials should be inversely correlated with the strength of ACC activity on the preceeding i trial. This is indeed what the authors found. This strengthens the hypothesis that one function of the ACC is the monitoring of conflict and the ensuing induction of control. This argument is supported by the observation that there was a trial-by-trial correlation with strength of ACC activity in the first incongruent trial with the reaction time on the following incongruent trial.

Another area which is frequently seen to increase activity with task difficulty is the dorsolateral prefrontal cortex (DLPFC). In many tasks the ACC as well as the DLPFC increase activity under the same conditions. It has thus remained unclear whether the two areas have different functions. Mainly based on human lesions, the prefrontal cortex has long been

implicated in cognitive control whereas ACC has been implicated in the inhibition of control. Also, DLPFC has been observed in the absence of ACC activity in working memory tasks (Fletcher et al., 1998), whereas ACC activity without DLPFC activity has been observed for incongruent response situations like in the Stroop task. To combine these two types of tasks, (MacDonald et al., 2000) have used a modified version of the Stroop color naming task. Before the start of a trial, an instruction was displayed as to whether the color of the word (ink) or the meaning of the word should be reported in this given trial. BOLD activity was examined both during the display of the instruction as well as during the response period. As previously reported, greater ACC activity was found for incongruent vs. congruent trials during the response period. DLPFC activity was elevated but equal for both types of trials. However, left DLPFC activity was different during the instruction period for color vs. word instructions. The authors conclude that this pattern of activation is indicative of a role for the DLPFC in cognitive control (setting of task) and the ACC in conflict monitoring. This finding is also supported by the study discussed in the previous paragraph (Kerns et al., 2004), which found that ACC activity predicted the extent of later PFC activity. That is, ACC signals a conflict, and due to this PFC takes the necessary actions ("control").

1.6.5 Error likelihood and reward

Another study (Brown and Braver, 2005) has challenged the above finding by proposing that ACC activity represents error likelihood rather than a conflict monitoring signal. The authors directly compared these two hypotheses using a go/nogo task with high- and low-error trials. Subjects were instructed to respond with a left or right button press to a left or right arrow (go

cue) appearing on the screen. However, in some trials, the go cue was reversed shortly after displaying it (no-go). The incidence of the no-go signal appearing was low in low-error trials and high in high-error trials. The subject was instructed by a color cue whether the current trial was a high- or low-error trial. Following the error monitoring hypothesis, activity should not differ between correct low- and high-error trials that were not followed by the no-go cue. In contrast, following the error likelihood hypothesis, activity in the high error-likelihood trials should be higher than the low error-likelihood trials, regardless of whether the no-go cue was actually displayed or not. This is because both trials have a higher likelihood of error, regardless of whether the error actually happened or not. Also, the comparison between high and low error-likelihood error trials that were not aborted (go cue) allows excluding effects of conflict (there is none) as well as error monitoring (there is none). The only parameter that differs is error likelihood. Indeed, the authors found that the fMRI signal measured in the ACC follows the error-likelihood hypothesis. That is, the signal was positively correlated with the potential negative reinforcement associated with a given trial. This is also in line with another study that found that fMRI ACC activity was positively correlated with the false alarm rate (Casey et al., 1997), because the higher false alarm rate was presumably related to trials where more errors could be made.

Another possible function of the ACC is the representation of some function of reward. This is in addition to uncertainty, as discussed in the previous paragraph. ACC modulation by reward is expected because it is known that the ACC has the highest density of innervation by the mesocortical dopamine system originating in the midbrain (Schultz and Dickinson, 2000). The midbrain dopaminergic system responds strongly to expectancy mismatches of reward. While

these dopaminergic connections itself are not capable of exciting neurons in the ACC (they are modulatory), dopamine is known to have a strong modulatory influence on synaptic plasticity and could thus influence ACC firing indirectly. Indeed, (Critchley et al., 2001) found that the extent of ACC activity (fMRI BOLD) is positively correlated with the amount of uncertainty (risk) as well as arousal (measured by skin conductance).

1.6.6 Lesions and human intracranial recordings

Given the contested nature of the function(s) of the ACC it is instructive to consult studies of the human ACC that look at causal rather than correlative effects. A rare opportunity of doing so are human patients that have lesions restricted to the ACC. One study investigated the performance in Stroop interference and go/nogo tasks in 4 patients with damage to the dorsal ACC and compared it with 12 healthy controls (Fellows and Farah, 2005). Overall reaction time in both the go/nogo task as well as the Stroop task were higher (slower) for the lesion patients. However, both the error rate as well as the size of the Stroop effect (percentage RT difference congruent vs. incongruent) were not different. Also, the modulation of the error rate and the Stroop effect size by high vs. low conflict was equal to controls. Slowing of the RT following an incongruent trial is considered one of the effects of cognitive control induced by ACC activity. However, it was observed at the same rate as in controls. Surprisingly, this study thus finds no difference in all measures of conflict that are traditionally considered functions of the ACC. Other studies (Stuss et al., 2001; Vendrell et al., 1995), however, agree with this finding. This finding is also in line with a number of non-human primate single-unit recordings that have generally failed to find neurons that respond to Stroop-like incongruence tasks. Rather, the neurons were found to

respond to functions of reward expectancy (Shidara and Richmond, 2002). The same has been found by a human study of selective cingulotomy patients (Williams et al., 2004) that allowed behavioral and electrophysiological measurements of performance on a task before and after resection. They found that performance was only impaired on trials that are related to reward, but not without.

The difficulty of studying ACC at the single-unit or LFP level is also illustrated by a recent study (Wang et al., 2005): the authors administered several different tasks to the same subjects and recorded LFPs from depth electrodes implanted for purposes of epileptic seizure localization. Tasks included auditory oddball detection, new/old word recognition (memory), and a reaction time task. Interestingly, the authors found activity that distinguishes between the different task elements (new/old, wrong/correct, rare/frequent) at each site—indicating that the function of the ACC at this particular site could not be attributed to one of these very different processes exclusively.