

Chapter 12

Functional and Structural MRI: Theoretical Background and Practical Aspects

Lukas Scheef and Henning Boecker

Abstract The following chapter will provide a treatise of MRI-based functional and structural imaging methods. It will start with a short summary of the physiological foundations of contemporary MRI-based functional imaging methods. We will continue with a general overview of pre- and post-processing steps that are relevant for the analysis of functional imaging time series. Data analysis issues will be presented as well as paradigm designs for assessing brain function. The subsequent discussion of structural image analysis methods will primarily focus on voxel-based morphometry (VBM), deformation-based morphometry (DBM), and diffusion tensor imaging (DTI).

12.1 Introduction

During the past 20 years, the field of functional magnetic resonance imaging (fMRI) has experienced a rapid development. At its beginning, the term fMRI was primarily referring to the acquisition and analysis of MRI time series that aimed to detect task-related changes in blood-oxygenation level dependent (BOLD) responses. Meanwhile, functional MRI-based applications have extended far beyond this initial narrow scope, now encompassing perfusion-based techniques (such as arterial spin labeling, ASL), as well as resting-state fMRI that observe spontaneous fluctuations of brain activity.

Moreover, the development of advanced scanner equipment and imaging sequences has afforded the time-efficient acquisition of structural imaging data: Today, almost every fMRI protocol contains a high-resolution T1-weighted sequence, and often, the scan protocol is also supplemented with a DTI sequence. Moreover, the development of (semi-) automated morphometric techniques for these structural imaging modalities

L. Scheef (✉) • H. Boecker
Functional Neuroimaging Group, Department of Radiology,
University of Bonn, Sigmund-Freud-Str. 25, 53105 Bonn, Germany
e-mail: Lukas.Scheef@ukb.uni-bonn.de; Henning.Boecker@ukb.uni-bonn.de

(e.g., VBM, DBM) has enabled unprecedented venues for analyzing the functional correlates of brain morphological features.

The fMRI field has been growing exponentially over the last 20 years, and due to the truly interdisciplinary character of research, it is almost impossible to become an expert on every aspect of this field. A neuroscientist most likely won't become an expert on network sciences, advanced statistics, or a mathematician. And it is not necessary to become one, as one does not have to become an MRI physicist to run fMRI experiments, but it is necessary to get a general overview and a good working knowledge in order to decide which path has to be followed to answer a certain research question. This is not only necessary in order to choose the optimal experimental design and data acquisition methods, but also in order to estimate the human and technical resources needed to conduct and analyze an experiment properly. This also includes aspects like boundary conditions for advanced analysis methods, which we plan to apply to the data, because they might impose restriction on paradigm designs, as well as pre- or post-processing strategies. This chapter aims to provide some general background and working knowledge that is needed to plan, conduct, and analyze fMRI experiments, and it also covers the morphometric analysis of structural MRI data.

12.2 Physiological Foundations of fMRI

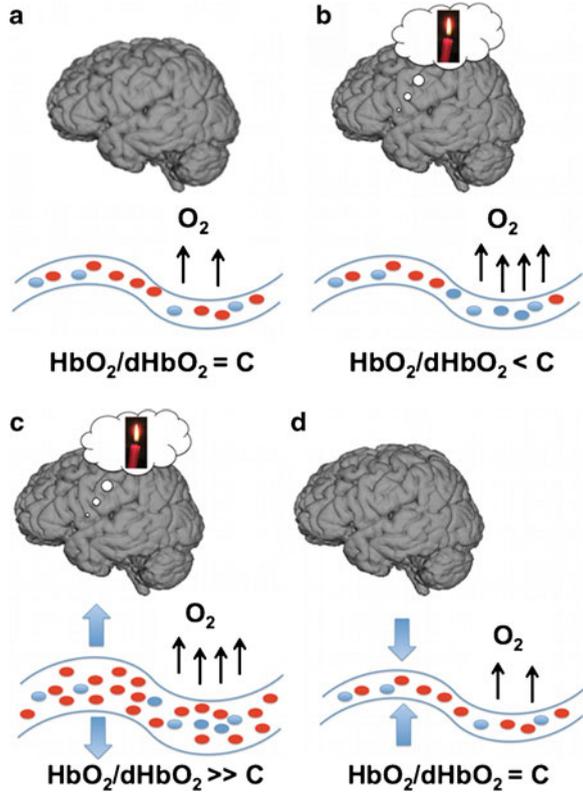
While several functional MRI techniques have developed in recent years, there is one aspect they have in common: They exploit hemodynamic responses that are triggered by changes in neural activation. When a focal brain activity occurs, the energy demand increases, and a plethora of physiological processes are triggered to ensure that sufficient energy is provided to support brain cell function (i.e., to keep the ATP concentration constant and the active ion channels up and running). In order to avoid any energy shortage, the vascular system reacts in this situation with a local vasodilatation, which leads to an increase of the local blood flow up to 40%. This temporary increase of regional cerebral blood flow (rCBF) in activated brain regions is a key process for all fMRI imaging techniques, and provides the physiological basis of *perfusion-based MRI* techniques (e.g., arterial spin labeling).

Yet, the short-term changes in rCBF have additional consequences (Fig. 12.1): In general, autoregulatory processes keep the global and local blood flow constant. The ratio C between the concentration of oxygenated and deoxygenated hemoglobin ($[\text{HbO}_2]/[\text{dHbO}_2]$) is also kept within constant margins during rest. However, when a certain brain region becomes (more) active, local oxygen consumption increases. Therefore, one observes a decrease of $[\text{HbO}_2]$, whereas $[\text{dHbO}_2]$ increases, and consequently, the ratio C between both decreases as well. Due to the drastic local flow increase in the activated brain area, much more oxygenated HbO_2 is actually delivered than consumed by the increased neuronal activity. This overcompensation leads to an increase of the ratio C , which exceeds the baseline level and remains elevated even after the local brain activation already went back to its initial level. This situation persists for up to 20–30 s before the whole system returns back to baseline values.

These short-lived changes in local blood oxygenation provide the physical basis for a second class of techniques: *BOLD fMRI*. As pointed out in Chap. 11, HbO_2 and

Fig. 12.1 BOLD Effect.

During “rest” the ratio between oxygenated and deoxygenated hemoglobin is fairly stable: $[\text{HbO}_2]/[\text{dHbO}_2]=C$ (a). When brain activity increases locally, this raises oxygen extraction, and therefore $[\text{HbO}_2]/[\text{dHbO}_2]<C$ (b). This is followed by vasodilation, leading to increased local blood flow and blood volume. As the increased blood flow delivers more oxygen than needed: $[\text{HbO}_2]/[\text{dHbO}_2]\gg C$ (c). After the locally increased brain activity stopped, the system slowly returns back to baseline again: $[\text{HbO}_2]/[\text{dHbO}_2]=C$ (d)



dHbO_2 have different magnetic properties. dHbO_2 is more paramagnetic and increases the local field, whereas HbO_2 is diamagnetic as the surrounding tissue (Pauling and Coryell 1936). These magnetic differences can have a profound effect on the MR signal of cerebral blood, which are mainly caused by dHbO_2 . The higher the dHbO_2 concentration in a blood vessel is, the more the local magnetic field is disturbed in the vessel and in the surrounding tissue. These magnetic field disturbances can be detected by using a T2-(Thulborn et al. 1982) or T2*-weighted sequence (Ogawa et al. 1990), because local field inhomogeneities cause a rapid dephasing of the transversal magnetization, which causes signal loss in these sequence types. In other words: If the dHbO_2 concentration increases, the local signal in T2- or T2*-sensitive sequences decreases, and vice versa—a phenomenon that is referred to as the BOLD effect.

12.2.1 Hemodynamic Response Function

As outlined above, local cerebral activity is followed by a local vasodilatation that causes an increased regional cerebral blood flow, which in turn induces local changes of blood oxygenation that are measurable via signal changes (i.e., BOLD responses) in T2- and T2*-weighted images. The signal changes that are induced by brain

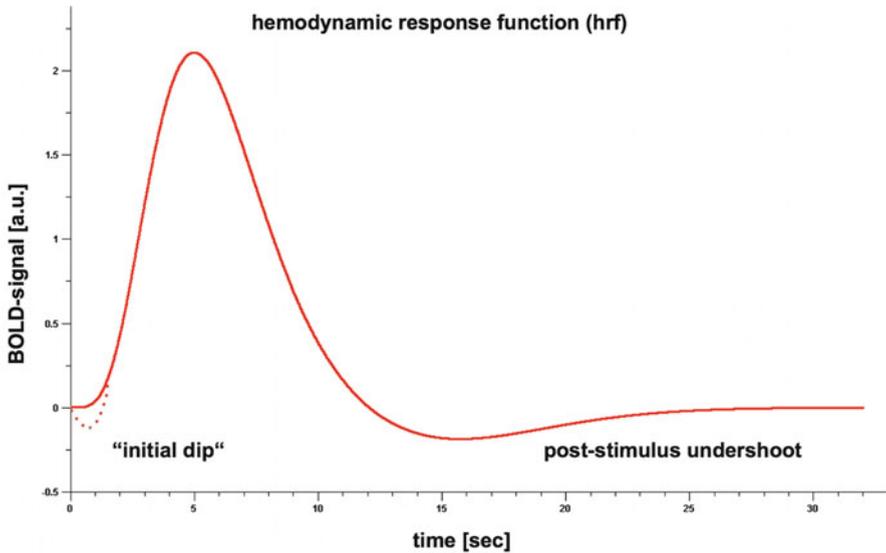


Fig. 12.2 Hemodynamic response function (*hrf*). Following a brief stimulus, $[\text{HbO}_2]$ decreases, which might lead to an initial signal decrease (initial dip). The main signal change is driven by the large-scale increase in cerebral blood flow and blood volume that increases the $[\text{HbO}_2]/[\text{dHbO}_2]$ -ratio and hence the signal intensity of the bold signal. Because the blood flow returns faster to baseline values than the regional blood volume one observes the poststimulus undershoot

activations (e.g., following a brief sensory stimulus) show specific temporal characteristics, which are usually referred to as *hemodynamic response function* (*hrf*; Fig. 12.2).

For a brief stimulus, one might observe an initial decrease of BOLD signal intensity within the first 1–2 s after the stimulus onset, which is often referred to as the “initial dip” (Ernst and Hennig 1994; Menon et al. 1995). It was suggested that the rapid increase of neuronal activity leads to a higher oxygen extraction fraction, which results in a decreased local $[\text{HbO}_2]/[\text{dHbO}_2]$ ratio, and therefore an initial BOLD signal decrease. Yet, the initial dip is very low in amplitude, and often hard to detect: Thus, its existence is still under debate (Behzadi and Liu 2006; Buxton 2001; Nirikko 2003; Vanzetta and Grinvald 2001; Yacoub et al. 2001b). However, recent data indicate that it might depend on the microscopic properties of the tissue and may therefore be not always and ubiquitously observable (Tian et al. 2010).

Meanwhile, the subsequent temporary hyperperfusion that is triggered by brain activation supplies the respective brain region with more oxygenated blood than is actually needed, resulting in a drastic increase of the local blood oxygenation. This change is reflected by a steady BOLD signal increase, which takes about 2–5 s after stimulus onset to reach its maximum. Thereafter, the BOLD signal slowly returns back to baseline (which takes up to 12–18 s), and even experiences a temporary decline below this level (the so-called poststimulus undershoot), lasting about 4–10 s. The poststimulus undershoot is thought to reflect the slow recovery of the cerebral

blood volume: In the so-called *Balloon model* (Buxton and Frank 1997; Buxton et al. 1998), which aims to provide a physiological model of the hrf, the shape of the hrf is described with the help of an input-state-output model, using the blood flow as input parameter, blood volume and the amount of deoxyhemoglobin as state variable, and the BOLD signal as output variable. The venous part of the vascular system is considered as a balloon, which is inflated by the arterial blood flow. In this model, the poststimulus undershoot results from a reduced clearance and dilution of the deoxygenated hemoglobin after the balloon has relaxed. The “balloon model” is simple and capable to explain the whole shape of the hrf, including the initial dip. However, it should be mentioned here that also alternative models exist, like the “windkessel model” (Mandeville et al. 1999) or the “steady-state model” (Hoge et al. 1999a, b).

12.2.2 *Properties of the hrf*

Even though the hemodynamic response to a brief stimulus can elicit rather long BOLD fMRI signal changes (which can last up to 30 s to return to baseline levels), it has three properties which are basically responsible for the success of this method in the context of brain activation studies: (1) It allows the detection of very short stimuli. (2) Its shape is fairly stable and therefore predictable across events, regions, and volunteers. (3) It is linear across a broad range of stimulus durations.

At the beginning of fMRI, especially the first point was a matter of debate. The early stimulation designs were conducted with rather long stimulus durations (around 30 s) that only elicited very small signal changes (about 2–5%). For this reason, it appeared very unlikely that much shorter stimulus durations would elicit any detectable signal. However, this conclusion proved to be false. Blamire and colleagues (1992) showed that a visual stimulus as short as 2 s was capable to induce a measureable signal change. Savoy et al. reduced the stimulus duration further down to 34 ms and was still able to detect an activation in the primary visual cortex (Savoy et al. 1994), while a motor activity of 500 ms duration was shown to be sufficient to elicit a BOLD response in the primary motor cortex (Bandettini et al. 1993). Within a stimulus range between 1 and 5 s the amplitude of the BOLD signal appeared to be linear (Dale and Buckner 1997), and it was suggested that the “responses to long-duration stimuli can be predicted from responses to shorter duration stimuli” (Boynton et al. 1996) because of its stability and reproducibility. Yet, the amplitude seems to be stable below 500 ms, and does not depend on the stimulus duration (Savoy et al. 1994): Hence, it cannot be considered being linear anymore. However, in general, the hrf is considered to be linear and time-invariant. In fact, the linear aspects explain up to 70% of the variance in fMRI data, depending on the exact mathematical shape, which is used to model the hrf. It might be modeled by a Poisson function (Friston et al. 1994b), gamma function (Boynton et al. 1996; Kruggel and von Cramon 1999), Gaussian function (Kruggel and von Cramon 1999), or a mix of Gaussian functions (Gossl et al. 2000). In practice, the implemented hrf

model is not actively chosen by the researcher, but hard-coded by the software package used to analyze the data. However, there are tools available to customize the hrf function (Woolrich et al. 2004b), which might be used to model it on an individual basis, to integrate pharmacological models into an fMRI analysis, to account for developmental changes or for the aging process, or for differences in the hemodynamic coupling itself, or to simply improve the model accuracy. In fact, the accuracy can be improved (up to 92%) if one considers the inter-subject variability of the human BOLD response. Aguirre, Zahran, and D'Esposito showed that shape of the hrf varies significantly between subjects, but was highly consistent within subject across scanning sessions if the same region is examined (Aguirre et al. 1998). Also significant variations can be observed between brain regions (Buckner et al. 1996; Miezin et al. 2000). So far, it remains unclear whether such differences are caused by vasculature differences, processing latencies, or by any factor influencing the hemodynamic coupling, as summarized by Muthukumaraswamy et al. (2011).

Having said that, one might wonder whether fMRI can be conducted robustly at all, but it has to be kept in mind that the influence of variability of the hrf is small compared to the share of variance that is successfully explained by (one of) the “standard” or “canonical” hrf functions. Those “unwanted effects” might even be useful to address very special questions like, for example, mental chronometry (Menon and Kim 1999 for review; Menon et al. 1998). However, for most designs it will be of less importance and can be ignored. The same holds true for the above-mentioned nonlinearity issues. They are mainly of relevance if the stimulus duration is short (<1 s) and/or the time between successive stimuli is below 2 s. On the other hand, these nonlinearities are of outermost importance if physiological models are investigated (Friston 2005; Valdes-Sosa et al. 2009), the neurovascular coupling is of interest (Magri et al. 2011), or causal models are used to explain fMRI data (Stephan and Friston 2010) or might be of interest for more advanced analysis methods like the investigation of graph theoretical models (Hartman et al. 2011).

12.2.3 Hemodynamic Coupling

As outlined above, local cerebral activity is followed by a vasodilatation leading to an increased regional cerebral blood flow and consecutively to (hopefully) measurable BOLD response. At first sight, this seems to be a very simple mechanism. But it raises the important questions: What exactly triggers the response, and what are the underlying cellular mechanisms? In other words: What does fMRI measure?

First of all, it must be emphasized that functional MRI techniques do not measure brain activity by itself, but capture hemodynamic response phenomena that are triggered by brain activation, which means that fMRI does not provide a direct measure for the metabolic demand caused by neuronal activity or of spiking activity. While the former notion holds true for both perfusion-based and BOLD fMRI, there are additional considerations that are specific for BOLD fMRI: Currently, it is assumed that the BOLD signal basically reflects the excitatory component of local

neuronal activity (Attwell and Iadecola 2002; Logothetis et al. 2001). However, transmitter release in the synaptic cleft itself can modify the local perfusion. So far, three mechanisms are known: (1) via nitric oxide triggered by glutamate release (Akgoren et al. 1996), (2) via astrocytes (Figley and Stroman 2011; for review), and (3) via neurotransmitters (noradrenalin, dopamine, serotonin (5-HT); see Harris et al. (2011) for review). A thorough treatise of these physiological mechanisms would be beyond the scope of this chapter, but there are helpful discussions in recent review papers (Attwell et al. 2010; Logothetis 2008; Mangia et al. 2009). Thus, differences in BOLD signal might not reflect differences in brain activity, which is a general problem of fMRI (Harris et al. 2011). And this is the crucial question when interpreting or better scrutinizing own results, because BOLD changes can be induced by every mechanism that influences the regional cerebral blood flow and oxygenation. Specifically, in the context of exercise-related studies, one has to keep in mind that exercise might interact with any mechanism on a cellular level, which might change the hemodynamic coupling and might mimic changes on a functional level.

However, this issue was raised since the first fMRI manuscripts were submitted for review and “shedding new light on the regulation of cerebral hemodynamics and metabolism by neural activity is still an ongoing effort and cherished goal in the fMRI community” (Kwong 2011), and should not hinder any researcher from using this methodology to answer research questions.

12.3 Paradigm Design and Optimization

There are several methodological issues that have to be considered while planning task-related fMRI experiments, which will be discussed in the following section. Beyond the general issue of appropriate task selection, the utilized MRI technique itself determines a couple of boundary conditions that constrain the experimental design. For example, there are technical requirements dictated by MRI data acquisition itself. Moreover, the signal properties of the hemodynamic responses that are measured with fMRI techniques impose certain constraints regarding the temporal characteristics of brain activations that should be met to detect these activations efficiently. Thus, the timing, duration, and presentation order of the different task conditions—usually referred to as experimental design—need to be optimized to enable a sufficient detection of the expected brain activations.

12.3.1 General Considerations

Because the amplitude of any induced signal changes is very small, the stimuli have to be repeated many times. The technical conditions that are present during the scanning of fMRI sequences add general constraints to the experimental design: For example, head motion could lead to image artifacts, which means that volunteer

movements need to be restricted as much as possible. As a consequence, responses are typically limited to simple button presses or tiny joystick movement. Spoken responses are difficult to record due to the noise emitted by the MRI system. Even though MRI-compatible microphones are commercially available, the mandibular motion causes image artifacts and should be avoided. Moreover, the scanner noise during image acquisition makes it is easier to deliver stimuli or commands visually than acoustically, even though it should be noted that dedicated headphones are available.

12.3.2 Selection of Appropriate Task Conditions

In a nutshell, fMRI experiments make inferences about the neural underpinnings of psychological processes by provoking psychological states under controlled conditions (usually by presenting standardized tasks) and by observing correlated changes in brain activity. Thus, the foremost challenge in conducting an fMRI study is the construction of experiments that are able to activate the psychological processes of interest. Usually, this is accomplished by comparing the brain activation level during an “active” task (that is assumed to activate the psychological process of interest) with the brain activation level during “baseline” or “control” conditions. The control condition might be either a low-level rest period (like fixation of a cross on screen), or an additional active condition. Optimally, the (active) control condition would differ from a condition of interest only regarding the key process of interest, whereas every other aspect is kept the same. For example, if one is interested to find a brain region that specifically activates when face stimuli are presented, one could present a reference condition with stimulus objects that have a similar shape, information content, image contrast, and colors as the target stimuli (e.g., scrambled faces), which are presented in the same frequency and for the same duration as the face stimuli. Even though not mandatory, an additional low-level baseline is helpful, because it resolves ambiguities. For example, the difference Faces (F)—NonFaces (NF) can yield positive results (i.e., stronger relative activation for faces) in the following constellations (where 0 refers to the low-level baseline condition):

1. $F > NF, F > 0, NF > 0$
2. $F > 0, NF = 0$
3. $F = 0, NF > 0$
4. $F < NF, F < 0, NF < 0$
5. $F = 0, NF < 0$

Noteworthy, only the second constellation reveals a truly face-specific activation. However, in order to decide whether one of the stimulus classes activated a region, one has to compare it against a (low level) baseline. The comparison with active control conditions aims to control for the possibility that at least some of the observed brain activations are triggered by other psychological processes that became co-activated during task completion.

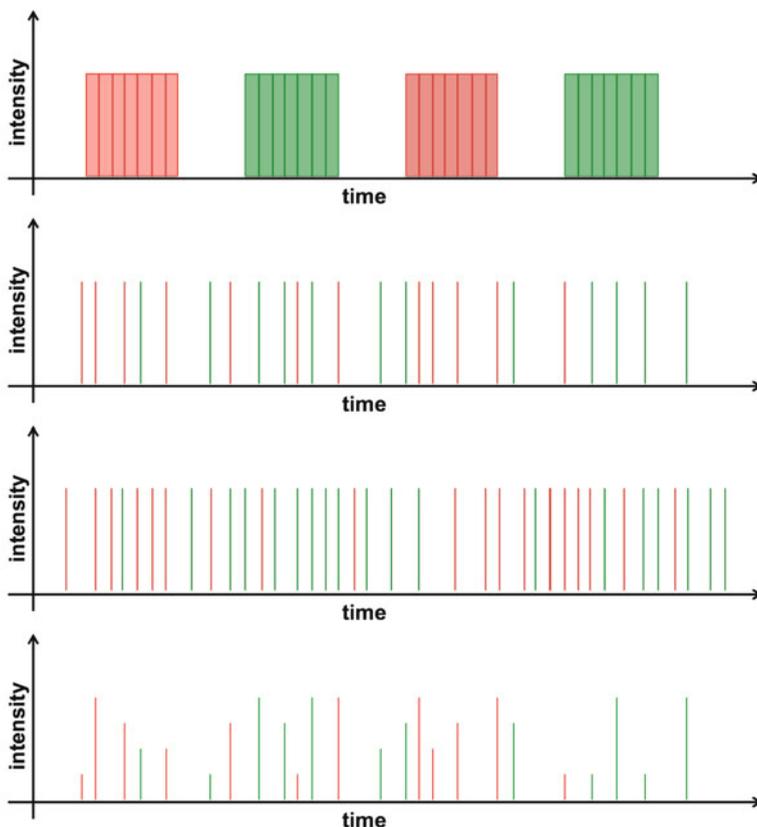


Fig. 12.3 *Paradigm design.* Four different paradigms are illustrated. From *top to bottom*: block design, randomized event-related design, mixed design, parametric design. Each paradigm consists of two event classes colored in red and green. In all cases, the low-level baseline is implicitly given by time periods without stimulus presentation (i.e., stimulus intensity equals zero)

12.3.3 Experimental Design Types

In general, task fMRI studies utilize two main classes of experimental design, the so-called block design and the event-related design (although there are also hybrid forms, so-called mixed designs, Fig. 12.3). Both design types differ regarding timing, order, and duration of stimulus presentation.

In a block design, stimuli of experimental and control conditions are concatenated into blocks of 10–30 s length, which are usually repeated about 4–6 times per condition. In this kind of design, the single stimuli presented within a block are not considered to be separate events, but are interpreted as a homogenous steady-state stimulation entity. Classical fMRI experiments to provoke motor

system (finger-opposition or finger-tapping task) or the visual system activation (“checkerboard stimulation”) are performed in this fashion.

The advantage of the block design is its simplicity and its power for detecting focal brain activation. As it will be discussed below, a block design leads to the highest signal changes and should be the first choice when designing an fMRI experiment. Since blocked designs aim to provoke a prolonged and continuous increase of fMRI signal intensity, they are most appropriate for investigating psychological processes that can be assumed to be continuously active during the task block. Moreover, block designs are most appropriate for experimental tasks where frequent switches between task conditions would disrupt attention and impair task performance (in fact, task switching may even cause an unintended activation of psychological control processes: see also Chap. 6). However, block designs have a couple of drawbacks, like habituation and predictability: The latter aspect is especially problematic for task paradigms that are critically dependent on the unpredictability of the trial order. For example, motor inhibition paradigms often present a larger proportion of go reaction trials which are randomly interspersed with a smaller number of inhibition trials in order to induce a prepotent go response tendency that has to be actively suppressed during the few inhibition trials (see also Chap. 6): Here, a massed presentation of inhibition trials makes their appearance predictable, and will probably reduce the activation of inhibitory processes. Moreover, the narrow spacing of single stimulus events at a fixed pace, which is typically found in block designs, will amalgamate the hemodynamic responses that are elicited by these singular events. This guarantees an extended continuous period of elevated fMRI signal, but the averaging character of this design type does not allow for extracting the hrf, or for detecting brain activity associated with events that are unpredictable in time, as, for example, the time point when the visual perception of a switch figure changes from one state to the other, like in a Rubin vase–face illusion image.

The second prototypical design type, the event-related design (ER-design), overcomes many of the latter obstacles. As the name already implies, such paradigms focus on brain activations elicited by single stimulus events, instead of fusing the responses of several stimuli, as it is found in block designs. Thus, event-related fMRI designs are especially appropriate for investigating phasic brain responses that are closely time-locked to specific stimulus events.

In their most basic form, event-related designs separate the presentation of successive events by long pauses of, say, 12 s or longer, so-called slow event-related designs. Due to these long intertrial intervals (ITI), the hrf that is elicited by individual stimuli can evolve and return to baseline before the next stimulus is presented, which means that there is no (or: few) overlap between the different hemodynamic responses. As such, slow event-related designs are especially useful for estimating the shape of the hrf.

Meanwhile, many event-related fMRI studies use rapid event-related designs. Here, much shorter ITI between successive stimuli are used than in slow event-related designs, increasing the subjective pace of the task. As a consequence, the hemodynamic responses elicited by the different events show a stronger overlap,

which makes it more difficult to identify the event-related signal peaks in the fMRI signal time series. To cope with this problem, rapid event-related designs vary the temporal distance between successive events, either by varying the ITI systematically (e.g., using a number of intervals in a range from, say, 1–5 s: “jittering”), or by adding so-called null events (i.e., by randomly interspersing blank trials with a fixed duration). The resulting variation in hrf overlap between successive trials can later be used in statistical analysis to isolate event-related signal changes.

Event-related paradigms allow for a maximum of design flexibility. For example, events for the different experimental conditions can be mixed in a random fashion, which means that habituation and predictability are minimized. Moreover, it is possible to perform a post hoc classification of the different events according to task performance (e.g., separate trials with correct and erroneous responses): A classical application of this strategy can be found in memory studies, where subjects can be scanned while they encode a series of items in an event-related design, and subsequently tested for their retention of this material. This way, it is possible to compare brain activation during the encoding of items which were subsequently remembered with the brain activation during the encoding of items that were subsequently forgotten, so-called subsequent memory effects (e.g., Brewer et al. 1998; Wagner et al. 1998). Another potential advantage of this design class is the opportunity to measure brain responses that are unpredictable in time, for example, the time point when the visual perception of a switch figure (e.g., a Rubin vase–face illusion image) changes from one state to the other.

Yet, the event-related design has the disadvantage of a reduced power to detect brain activation as compared with a block design, which means that longer experimental durations are necessary to achieve comparable results. This will become especially salient in slow event-related designs, but is also true for fast event-related designs. While the stimuli of different classes can be as close together in time as 2 s and can be still resolved, because the hrf is linear in this range (Dale and Buckner 1997), successive trials should not be spaced closer in time.

Both of the above-mentioned design types can be combined in a hybrid format, so-called mixed designs. Here, different experimental conditions are presented in a block-wise fashion, but the temporal spacing of single events within a task block is systematically varied, or different event classes are intermixed, which allows for a separation of event-specific responses in later statistical analysis. This approach is especially useful if researchers are interested in segregating tonic and phasic brain responses. For example, Donaldson et al. (2001) used a mixed design approach for a recognition memory task where subjects had to recognize items that they had learned before. The design enabled to differentiate between activations reflecting the “retrieval mode” (i.e., the attempt to remember the stimuli, which was tonically present across the whole task block), and phasic responses, which reflected the successful recognition of remembered items (which were time-locked to the singular events that were recognized successfully).

So far, the different design types differ regarding the timing of the stimuli. The conditions operationalize the independent variable of interest (i.e., a psychological process) in a categorical fashion, that is, as fixed factors, where task conditions are

treated as discrete on-off states (i.e., are either present, or absent). Yet, it is also possible to treat the independent variable as a continuous variable. If the stimulus intensity for a task condition (i.e., pain intensity, cognitive demand, and finger-tapping frequency) is modulated throughout an experiment, the design is called parametric design. This does allow for measuring not only linear relationships (i.e., the stronger the stimulus intensity, the stronger the brain activation), but also exponential or curvilinear relationships. This design pattern is applicable to both block and event-related designs.

A good review on fMRI design types is given by Amaro and Barker (2006) or Chein and Schneider (2003). Regarding potential pitfalls, the reader is also referred to the discussion by Savoy (2005).

12.3.4 *Optimizing Paradigms*

When designing an fMRI experiment, the first and most important decision to be made is the choice of the paradigm design type. As a general rule, it can be stated that a blocked design should be preferred over an event-related approach, because it provides the optimal statistical power to detect BOLD changes. However, it does not allow for an investigation of the hrf, the analysis of single stimuli, and might not be adequate for the underlying neuroscientific question, which is of course the most crucial aspect. The boundary conditions for the latter point are (1) a condition is represented by only two states, which can be switched on and off (for example: self paced finger tapping: yes or no), (2) there is no need to distinguish between single stimuli within a block, and (3) processes cannot be separated from another process by a couple of seconds. An in-depth discussion can be found in Zarahn et al. (1997).

If the shape of the hrf is known (Aguirre et al. 1998), the optimal duration of a block is 14–20 s (Zarahn et al. 1997); otherwise a duration between 30 and 40 s might be used. A block is usually repeated 4–8 times. Due to slow scanner drifts and low-frequency noise (with its 1/f characteristic), it is not advisable to use larger block durations than 30–40 s (Zarahn et al. 1997). In other words, longer block durations increase the risk of misinterpreting these non-task fluctuations in BOLD signal as task-induced changes in brain activation. Therefore, BOLD fMRI are problematic if researchers are interested in examining psychological processes that cannot be switched on and off within a few seconds (e.g., mood or pain states). However, in these situations perfusion-based fMRI (see below) may provide a useful alternative, since these techniques are less sensitive to these forms of low-frequency noise. Considering the stimulus presentation within a given block, the number of stimuli presented within a single block, and the spacing between these stimulus onsets (stimulus onset asynchrony, SOA) primarily depends on the nature of the experiment and the used stimuli. In most block designs, the conditions have a fixed order, in which the conditions alternate, and are interleaved by a low-level rest condition. Considering the ordering of the different task conditions, there are a variety

of different approaches, for example the classical “boxcar design” (where an active (A) and a control (C) condition are presented in alternation, e.g., A-C-A-C-A-C), or “castle designs” (where an active (A) and an active control (C) condition alternate with a low-level resting baseline (R), e.g., R-A-R-C-R-A-R-C-R). The advantages and pitfalls of these different variants are extensively reviewed in Chein and Schneider (2003).

Event-related designs provide the maximum degrees of freedom for the paradigm design, but are more difficult to optimize. The simplest design is built by events of only one category separated far enough to allow the hrf to return (almost fully) to baseline (i.e., in a slow event-related design). In this situation, the most sensitive setup (=statistical power/time) is reached if the time between two successive stimuli, the interstimulus interval (ISI), is about 16 s. If the scan time is not crucial, the sensitivity could be further increased if the ISI becomes larger than 20–30 s. If more than one stimulus categories are used, the power for intercategory differences can be increased by reducing the ISI well below 16 s as long as they are in randomized order. However, the sensitivity for detecting the evoked response of a single category necessarily decreases. If the ISI is not fixed but “jittered” around a mean ISI, the sensitivity/time can be increased, but will be always below that of a block design. Designs with long ISIs can also be used to extract the shape of the hrf. Due to the small within-subject and within-trial variability of the hrf (Aguirre et al. 1998), this approach can be used to detect even small differences in the onset of neuronal activity between conditions, as long as stimulus order, ISI, and jitter are fully randomized. If the number of conditions increases ($n > 2$), or the contrasts become more complex, it is almost impossible to optimize these parameters manually. However, algorithms and free software tools are available that help to generate stimulus sequences and timing parameters that are optimized to detect contrasts of interest (Wager and Nichols 2003).

The number of events per condition that are needed to gain a sufficient signal is difficult to determine in advance. If the sampling rate of the fMRI scanning sequence is sufficiently high, it might detect even a single event, while if a high-frequency rapid event-related design is chosen, even 50 events or more can be insufficient. As a rule of thumb, at least about 30–40 events per condition should be presented. Yet, the necessary number of events per condition depends on the design and the underlying question.

The same holds true for the number of different conditions: If the number of conditions becomes too large, the balancing and the optimization become increasingly difficult. Moreover, the experiment is simply at risk of becoming too long. From a practical point of view, a single fMRI run that lasts longer than, say, 30 min means a very high burden for a volunteer, not only due to the noise of the MR and the narrowness of the magnetic bore, but also due to the attentional demands of the task which may become increasingly fatiguing.

This section was only able to give a broad overview of the topic fMRI design optimization. For further details, we like to refer to more extensive treatises (Dale and Buckner 1997; Friston et al. 1996; Friston et al. 1999; Josephs and Henson 1999; Zarahn et al. 1997).

12.4 Acquisition of fMRI Time Series

12.4.1 Field Strength

While experimental design is a crucial issue for fMRI studies, adequate scanner equipment is another relevant topic, because the technical limits of these systems necessarily will often restrict the signal-to-noise ratio (SNR) for measuring the physiological signal changes that are evoked in an fMRI experiment.

The most expensive way to maximize the signal-to-noise ratio is to maximize the field strength. As noted above, the BOLD effect is a compound of both T2 and T2* effects, which have a distinct physical basis. T2 is more sensitive to the intravascular effects, whereas T2* is more relevant for the extravascular tissue. Ogawa was able to show that the effect of intravascular component scales linearly with the field strength, whereas the tissue component scales exponentially (Ogawa et al. 1993) and shifts therefore from the vessel compartment more and more to the surrounding tissue (Gati et al. 1997).

Unfortunately, field strength is not an upgrade option, and the system prices scale by a factor of about 1.000.000 US\$/T. Meanwhile, higher field strengths also increase the risk of image artifacts, especially outside the brain. Therefore, it has to be critically discussed with all users of the system whether the higher price of a higher field strength system really pays off. Currently, 3T can be considered as the standard field strength for MRI systems used for humans, but it has to be pointed out that most fMRI experiments can also be performed sufficiently with 1.5T scanners. The MR physics at these field strengths is fully understood and the vendors solved most technical problems. They are highly optimized and the leeway for further improvements is limited.

On the other hand, ultra-high-field scanners (i.e., 7T or even higher) should primarily be considered as being experimental systems. They show a drastic increase in BOLD sensitivity, but they are still far from being optimized, and a dedicated staff of MR physicists and technicians should be (fully) available in order to keep the system up and running (including sequence optimization). An in-depth discussion of different aspects of high-field- and ultra-high-field fMRI is given elsewhere (Gati et al. 1997; Hoenig et al. 2005; Triantafyllou et al. 2005; van der Zwaag et al. 2009; Yacoub et al. 2001a).

12.4.2 MR Sequences

There are mainly two criteria for MRI sequences that have to be met for BOLD-based fMRI: They have (1) to be sensitive to T2 or T2* changes, and (2) to allow fast data acquisition, since the time course of BOLD signal changes needs to be

observed with reasonable temporal resolution. Both requirements are met for so-called Single-Shot Gradient Echo Planar Imaging sequences (sshGE-EPI), when $TE \sim T2^*$ for gray matter (see also Chap. 11). They are the “working horse” for fMRI. With this sequence type, it is possible to acquire a multi-slice scan volume that covers the entire brain, while maintaining an acceptable spatial resolution about 3 mm voxel size, within a time frame well below 3 s (usually referred to as Time of Repetition, TR). Since a scan volume is acquired slice by slice, this means that the BOLD signal measurement for each single voxel is repeated every TR. Depending on the goals of the study, researchers may wish to achieve a higher sampling rate. For this purpose, there are different strategies available: For example, we can increase imaging speed simply by reducing the number of slices acquired. However, this is necessarily paid off with reduced brain coverage. Another option to accelerate image acquisition is the use of half-Fourier imaging or parallel imaging (see Chap. 11). In combination with another sequence type, which is called PRESTO (PRinciple of Echo Shifting with a Train of Observations (Liu et al. 1993)), a whole brain coverage can be accomplished within 500 ms and an isotropic spatial resolution of 4 mm (Neggers et al. 2008). However, a slight disadvantage of the latter might be the drastic increased noise level accompanied with this sequence.

While optimizing an fMRI sequence, one has to bear in mind that a high sensitivity for the BOLD effect, i.e., $T2^*$ -sensitive sequences, also results in a high sensitivity for static susceptibility gradients, which are especially present at the skull base, or close to the mastoid process. These gradients lead to a signal void and huge image distortions in adjacent brain regions. Using standard sequences and setups for whole brain image acquisition, it is therefore almost impossible to reliably scan temporobasal structures, like the amygdala or the hippocampus, or the orbitofrontal cortex. However, there is a growing literature discussing methodological approaches that can help to reduce these effects, for example by increasing the spatial resolution, using oblique image acquisition planes (Weiskopf et al. 2007), and parallel image acquisition schemas (Schmiedeskamp et al. 2010) which reduce these artifacts to an acceptable degree. Homogenizing the local magnetic field by applying additional magnetic field gradients and preparation pulses (Deichmann et al. 2003) or the use of intra-oral shims (Osterbauer et al. 2006; Wilson and Jezzard 2003) can further improve the image quality.

If long experimental conditions are needed (e.g., blocks longer than 30 s), classical fMRI sequences may not be ideal solutions due to scanner drifts and the low-frequency noise profile. Here, arterial spin labeling should be considered (Aguirre et al. 2002; Detre and Wang 2002; Wang et al. 2003). ASL is also useful when investigating states, which cannot easily be switched on or off, such as tinnitus. ASL time series consists of a series of successively acquired image pairs (unlabeled+labeled). Therefore, they are less susceptible to scanner drifts and low-frequency noise, and can be considered as method of choice in such situations. However, it should be mentioned that ASL sequences are difficult to implement and demand exquisite scanner stability.

12.5 Preprocessing of fMRI Time Series

Before fMRI data can be analyzed statistically, there are a number of preprocessing steps that have to be performed. First of all, the data have to be corrected for temporal differences between image slices that are due to the sequential image acquisition character of single-shot sequences (*slice-timing correction*). Moreover, image preprocessing needs to remove head movement effects (*realignment*). The spatially low-resolution functional imaging time series should undergo a *coregistration* to a high-resolution anatomical scan, allowing for later anatomical assignment of any imaging result. For group analysis, the imaging data have to be spatially transformed into a standard anatomical space (*spatial normalization*) and finally undergo a *spatial smoothing*, which compensates for residual individual anatomic differences, and increases the SNR for subsequent statistical inference. These preprocessing steps will be explained in the following subsections:

12.5.1 Slice-Timing Correction

The standard fMRI sequence is a single-shot EPI, which means that an image volume is acquired slice by slice, which takes about 3 s, depending on echo time, spatial resolution, and number of slices acquired. In other words, there is a huge time gap between the first and the last slice acquired within a single volume. This means that the BOLD signals that are collected from voxels in different slices reflect the state of the brain at slightly different points in time (i.e., have a different slice timing). For a block design, such a time gap is usually of less importance, because the design aims to average across the blocks of the same condition. Event-related designs, however, are designed to analyze the time series on the basis of individual events, which means that they critically rely on an exact timing of the BOLD responses. Therefore, slice-timing differences should be corrected. For short TRs (<2 s), the small timing differences can be accounted for by adapting the model setup to analyze the data by introducing the first derivative of the hrf time course into the statistical model (see below). A disadvantage of this approach is a reduction of statistical power, which is accompanied with this method, and the limitation to the small TRs.

However, because of the temporal characteristic of the HRF (i.e., because it is slowly varying in time), this problem can easily be solved by interpolation. The simplest approach is linear interpolation: When the image intensity in a voxel is known at time point t_1 and t_2 , the intensity within each time point in-between can be estimated by drawing a line through both time points, determining its equation, and calculating the image intensity at $t_1 + \Delta t$. To minimize the time difference Δt that has to correct for, it is usually advisable to use the time point of the middle slice as a reference for the slice-timing correction. By doing this, the correction is limited to a maximum of $\frac{1}{2}$ TR. Alternative methods are spline- or sinc-interpolation, which

differ simply with regard to the number of reference points needed for the interpolation process, and the fitted function. See Sladky et al. (2011) for a recent discussion of this topic.

It might be argued that in the presence of head motion (which more or less always exists in fMRI time series), this approach might be not applicable, because the method relies on a rigid model (= the voxel coordinate does not change between two time points), which is not valid in the presence of motion. Others suggest realigning (i.e., motion correction) first, before the slice timing is performed, which might violate the assumption that two voxels with the same spatial coordinate have the same time lag relative to the reference slice. However, both approaches rely on the assumption that the differences (either in time or in space) between two consecutive volumes are small and, hence, both methods are suboptimal. From a mathematical point of view, the best solution for this problem would be to combine both approaches in one algorithm. This was recently accomplished and, indeed, showed a more accurate image reconstruction (Roche 2011).

12.5.2 Motion Correction

The motion correction procedure corrects for head motion during and across fMRI sessions, which can result in slightly different spatial positions of the brain between adjacent image volumes. Since the procedure acts on the image data that are acquired with exactly the same image modality, the motion between successive acquired image volumes can be corrected using a rigid body transformation (with three rotations and three translations). The optimal transformation can be determined using a minimal least square algorithm, which minimizes the residual sum of squares of the difference between two image volumes (Ashburner and Friston 2007b). At the end of this procedure, the whole time series is realigned either to the first volume, or to the mean of the time series. The resulting motion correction parameter can be later used as covariates of no interest in the statistical analysis step: This helps to account for residual noise associated with head motion, which even after realignment is typically large compared to experimentally induced variance. Another option that reduces this source of variance is to use not solely a rigid body transformation but to add a deformation-based term. The rationale for introducing a deformation term is based on the fact that magnetic field is not homogeneous in the field. Consequently, T2* images are a bit distorted. The distortion depends on the orientation of the head in the external magnetic field, and hence, head motion leads to variable deformations of the EPI images along time. This can be corrected by estimating the field changes and corrected by applying appropriate deformations during the motion correction procedure (Andersson et al. 2001). This method, which is usually referred to as “unwarping”, provides a better reduction of the motion-induced (and therefore unwanted) residual variance.

12.5.3 Coregistration

The price that is paid for the fast image acquisition of fMRI time series is the low spatial resolution. For example, a standard voxel size at 3 T is about $3 \times 3 \times 3 \text{ mm}^3$, which means that the intrinsic spatial resolution of the functional images is in general not sufficient for proper anatomical assignment of detected activation clusters. This problem can be solved by a coregistration of the fMRI time series onto a high-resolution anatomical scan of the participant. Typically, this is a T1-weighted dataset with a spatial resolution of about $1 \times 1 \times 1 \text{ mm}^3$. Because the image contrast and the spatial resolution of both modalities are different, the rigid body algorithm mentioned above cannot be used. The most common algorithm to solve this problem is called “mutual information maximization”. Mutual information is a measure of statistical dependency between two variables. It is zero if two variables are not associated with each other and maximized otherwise. When using the image intensity histograms of both modalities as variables, this method can be used to find the optimum transformation to overlay the one image modality on the other. Because histograms are used (i.e., counting the number of occurrences of voxel intensities within certain intervals), this method can be used independently of the spatial resolution of the compared modalities (Woods et al. 1993).

12.5.4 Spatial Normalization

The term spatial normalization refers to a spatial transformation procedure that transforms a brain dataset into a common frame of reference. Spatial normalization serves two purposes: (1) the compensation of individual anatomical differences between subjects, which is relevant for group analyses, and (2) the simplification of anatomical orientation, when finally assigning activation foci to specific anatomical structures (Brett et al. 2002). The latter aspect also simplifies the comparison with research results published in the literature, even across the imaging modalities. This is achieved by spatially transforming the data into a common space of reference. The most popular frame of reference is provided by the atlas from Talairach and Tournoux (TTA) (Talairach and Tournoux 1988), which was based on detailed anatomical dissection of the 60-year-old women. Usually referred to as “Talairach space”, the authors proposed a coordinate system where the origin is located in the anterior commissure, and the brain is oriented in a way that the x -axis runs through the posterior commissure. With appropriate scaling applied, the atlas can be used as a lookup table, assigning coordinates to major cortex structures, including Brodmann areas. In the last years, an atlas developed by Montreal Neurological Institute (MNI) largely replaced the TTA. It was created by averaging 152 high-resolution T1-weighted MRI datasets, using the same origin and axis orientation as defined by the TTA. The averaged datasets serve as a template, which can be used for coregistration of individual datasets to a common reference space. A large number of normalization algorithms and reference atlases have been published for the normalization process (Ashburner and Friston 2007a).

In a recent publication, different algorithms and atlases were compared against each other (Klein et al. 2009) and showed stable and reliable results irrespective of study populations and utilized algorithms.

A completely different approach is used by Anders M. Dale, Bruce Fischl, and colleagues, which does not aim to optimize the superposition of complete datasets onto a 3D template but to perform a cortical surface-based inter-subject alignment (Dale et al. 1999; Fischl et al. 1999a, b). This method is highly accurate, but restricted to the cortical surface, that is, does not cover subcortical regions or cerebellum. For comparison of the different spatial normalization procedures, we also refer to Gholipour et al. (2007).

12.5.5 Spatial Smoothing

Spatial smoothing refers to a procedure where signal intensities are averaged across neighboring voxels. Hereby, the voxels far away contribute less to the average, as compared to the direct neighbors of a voxel. The weighting function is usually a Gaussian bell curve, called Gaussian smoothing kernel, which is specified by its full width at half maximum (FWHM).

The advantage of the smoothing process is fourfold. First of all, it further increases the anatomical reliability of group studies, because the activation clusters are spread out in space, which increases the probability of overlap across subjects. Secondly, smoothing increases the signal-to-noise ratio, as every averaging process reduces noise components in the data. Thirdly, regarding the analysis of time series, the smoothing process has the advantage that it “normalizes” the noise structure of the BOLD signal (central limit theorem). The fourth rationale for smoothing lies in the statistical inference procedure. It reduces the number of independent comparisons within a dataset, and allows for less conservative statistical thresholding. As a rule of thumb, the smoothing kernel is chosen to be twice to three times larger than the spatial resolution of the dataset. However, if small structures are primarily of interest (e.g., amygdala, hippocampus, or brain stem structures), smaller smoothing kernels should be chosen (because the “matched filter theorem” dictates that the width of smoothing kernel should match the spatial dimensions of the structure of interest).

12.6 Analysis of fMRI Time Series

The easiest way to analyze fMRI time series is to average the images for the different conditions, and to perform a voxel-wise *t*-test to find all those voxels whose signal intensity (or brain activity level) differs significantly between the conditions. In this analysis strategy, the shape of the hrf is not taken into account, and it is difficult to extend this method to analyze event-related designs. An alternative approach is to model the *expected* time course of a condition and calculate

voxel-wise cross-correlation coefficients between the expected time course and the model time course (Bandettini et al. 1992). The resulting cross-correlation map provides information which voxels show condition-dependent BOLD changes. Even though these maps can be calculated very easily and can be applied to block and event-related designs, this approach has the disadvantage that it is difficult to directly compare two conditions with each other. Moreover, the cross-correlation takes only the shape into account, but ignores the amplitude of the respective activation.

12.6.1 The General Linear Model

The most flexible approach, which can be considered the current “gold standard” for fMRI analyses, is the use of the general linear model, which was initially proposed by Karl Friston and colleagues (Friston et al. 1994a; key concept). The main idea behind this approach is to model the time course for each condition in each voxel, and to estimate the response magnitudes by minimizing the residual error. Having done this, the differential effect of two conditions can be calculated by comparing the two amplitudes. The significance of the difference is estimated by comparing the amplitude difference with the residual error. To illustrate the procedure, let’s consider the following *gedankenexperiment* which is outlined in Fig. 12.4: Imagine a block design paradigm with two active conditions (A and B), which are repeated eight times in alternating order. A low-level baseline separates both conditions. The block duration is 15 s. 160 volumes are acquired (TR=3 s). The time course for each condition can be modeled by setting the value of the model function $X_{A/B}(t)=1$ for those time points when the corresponding condition is “active” and $X_{A/B}(t)=0$ for all other time points. In this case, we do not need to model the baseline condition explicitly. Now we can try to express the observed time course $Y_n(t)$ for any given voxel as a linear superposition of both model time courses plus an error term $\varepsilon(t)$:

$$Y_n(t) = \beta_{A,n} X_A(t) + \beta_{B,n} X_B(t) + \varepsilon_n(t). \quad (12.1)$$

Hereby, epsilon is assumed to follow a normal distribution. β represents the amplitude and has to be estimated for each condition and voxel separately. β can be estimated by minimizing the sum of squared errors, that is, by minimizing the term:

$$\sum \varepsilon_n^2 = \sum (Y_n - (\beta_{A,n} X_A + \beta_{B,n} X_B))^2 = \min! \quad (12.2)$$

Equation (12.1) can be rewritten using matrix notation:

$$Y = X\beta + \varepsilon. \quad (12.3)$$

Because X contains the experimental design, or in statistical terms the “predictors”, it is called design matrix. Y represents the real measured data or the response, whereas β are the unknown coefficient predictors and ε the error term. Figure 12.4

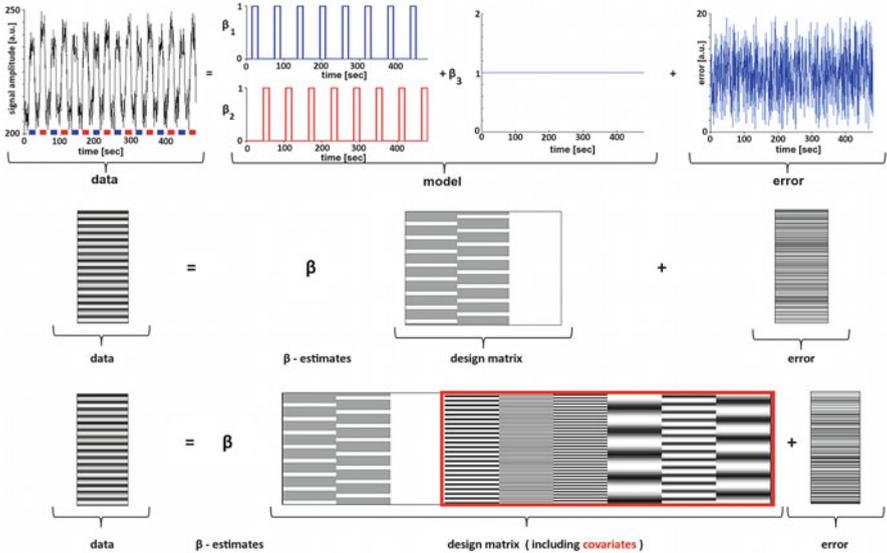


Fig. 12.4 *General Linear Model (GLM)*. The main idea behind the GLM is to explain as much variance of the measured data as possible by fitting a model to the observed data (*top row*). Each condition (Condition A: *blue*; Condition B: *red*) is modeled by a separate regressor. Both conditions superimpose linearly onto the mean signal intensity, which is represented by the third term. The optimal model parameters, the beta estimates, are found when the residual error term is minimal (*top row*). The model can be rewritten by concatenating the conditions into a so-called design matrix. Each column represents a single regressor. Time is coded along the row. Its graphical representation is shown in the middle row. The major advantage of the general linear model is its flexibility. Additional conditions can be added easily by adding an additional regressor. Moreover, covariates of no interest can be added as outlined in the *bottom row*

shows a graphical representation of the design matrix of the described experiment. The design matrix decodes each time point in a different line, whereas each condition is coded column-wise. The beauty of the General Linear Model (GLM) is its simplicity and extensibility. One can easily add covariates by just adding an additional term. A linear term (over time) might be added to compensate for linear drifts, or the motion parameters, derived from the motion correction procedure, to correct for head motion during the experiment:

$$Y_n(t) = \beta_{A,n} X_A(t) + \beta_{B,n} X_B(t) + \beta_{drift,n} + \beta_{motion,n} X_{motion}(t) + \varepsilon_n(t). \quad (12.4)$$

However, using the GLM as described here presumes a couple of assumptions: (1) the same model time course can be used for all regions activated by a certain condition, (2) all voxels are independent in time and space, and (3) the error term is normally distributed (mean=0; variance σ^2), and its distribution is constant.

12.6.2 Detecting Brain Activations

So far, it was described how the data can be modeled using the GLM. However, the most crucial point is not answered yet, that is, how one can detect those regions activated by a task or which regions are activated stronger by task A compared to task B . Using the GLM, this questions can be properly addressed by comparing the estimated parameters $\beta_A - \beta_B$. If the difference is large compared with the unexplained variance, this would usually be considered significant. This so-called effect can be calculated formally by multiplying β with a contrast vector C . This vector depicts the β -estimates of interest. In the given example (12.4), the question which voxels are more activated under condition A as compared to condition B could be accessed by the $C = [1 \ -1 \ 0 \ 0]$. The multiplication $C * \beta$ yields: $\beta_A - \beta_B$. In order to address $B > 0$ (i.e., condition B activated above baseline), C would be $C = [0 \ 1 \ 0 \ 0]$, or $C = [0 \ 0 \ 1 \ 0]$ for the linear drift term accordingly. It should be noted here that $B > 0$ could also be expressed as differential contrast if the rest condition is explicitly modeled.

The statistical significance of any effect can be tested either using T- or F-statistics (generally implemented in all fMRI software packages; see below), which will depend on the specific contrast in question: Both tests compare the effect explained by the contrast with the residual variance and calculate a corresponding T- or F-value. Because this procedure is performed for every voxel of the brain, a corresponding T- or F-value is obtained for every voxel yielding a statistical three-dimensional map, which is thresholded depending on a chosen significance level. For better localization of the activated foci, the result can be overlaid as a color-coded map onto a high-resolution structural dataset (Fig. 12.5).

12.6.3 The Multiple Comparison Problem

One critical point for voxel-wise statistical analyses of imaging data is the large number of comparisons that are performed. A standard fMRI volume consists of $64 \times 64 = 4,096$ voxels/image, or 122,880 voxels if 30 slices were acquired. Each voxel is tested against the null hypothesis (i.e., that it is not activated), or that there is not a difference between two conditions, respectively. A significance level $p < 0.05$ means that if a difference is observed the probability of observing such a difference only by chance (= false positive) is less than 5%. However, it also means if 100,000 tests are performed, one could expect 5,000 false positives just by chance! The classical correction for this so-called multiple comparison problem is the Sidak correction or its approximation for large N , the Bonferroni correction, which divides the desired significance level by the number of independent tests performed. Even though the Bonferroni correction is easily calculated, it has a couple of disadvantages. If the significance level is chosen to be 0.05 and 100,000 voxels are tested, the Bonferroni corrected significance level would be $0.05/100,000 = 0.0000005$. Thus, it is very conservative from a practical point of view. To avoid making too many

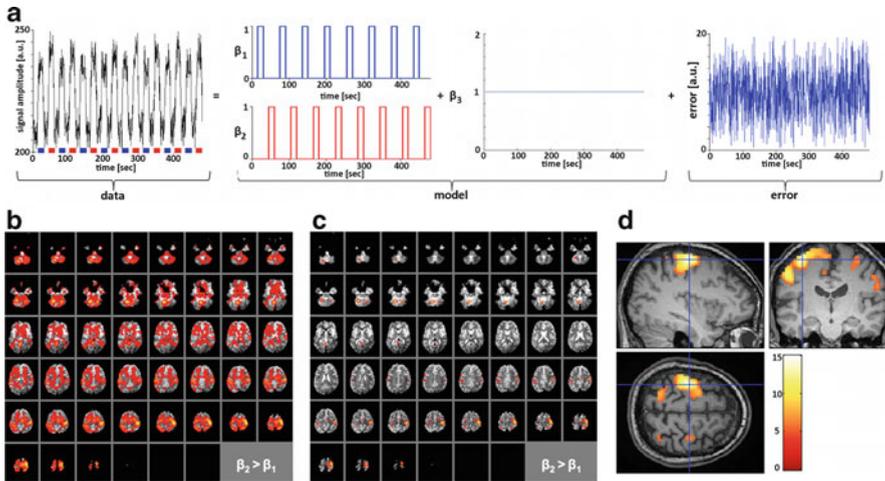


Fig. 12.5 *Statistical mapping.* Based on the general linear model (a), the parameters β_1 , β_2 , and β_3 are estimated. In order to access if and where condition A (blue) reveals a larger amplitude compared with condition B (red), the difference $\beta_1 - \beta_2$ is calculated and compared with the residual error term. This leads to a color-coded statistical map showing all voxels, where $\beta_1 - \beta_2 > 0$. The map is superimposed onto the mean-fMRI volume, which was constructed by averaging the whole time series after the realignment procedure. After applying an appropriate statistical threshold (in this case $p < 0.05$, FEW corrected) only the significant voxels are left over (c). By overlaying the data onto a high-resolution isotropic structural dataset, the anatomical assignment of the activated voxels can be improved (d). In the presented example, the subject performed a finger-tapping task. After appropriate thresholding, the sensory motor system is nicely depicted. The color-coding scheme for the t -values as given in (d) was used for all panels (b–d)

false-positive decisions, we risk missing true positives. Moreover, the Sidak or Bonferroni correction is not valid from a mathematical point of view, since neighboring voxels are not independent from each other, because activated regions are (usually) not confined to one voxel. Neighboring regions are heavily interconnected and supplied by the same vasculature, which means that it has to be assumed that the oxygenation level is similar in the close neighborhood. If one could calculate an effective number of “independent” voxels, the Bonferroni correction could be applied using a smaller N . The problem here is how to determine the number of independent voxels.

To some extent, this problem can be reduced by a spatial smoothing of the data, which was already discussed in Sect. 12.5.5: Worsley and colleagues showed that by smoothing the data with a Gaussian kernel, the number of independent comparisons can be estimated by replacing the number n of voxels by a number R of so-called RESELS (= resolution elements):

$$R = V / \text{FWHM}^3, \quad (12.5)$$

V = volume to be analyzed, FWHM = FWHM of the Gaussian smoothing kernel (see Worsley et al. 1996 for details). As a rule of thumb, using smoothing kernels

with FWHM of $3\times$ the voxel size reduces the multiple comparison problem by a factor about 30. This description is surely oversimplified, but should give an impression of the idea behind it. The mathematical theory applied here is called Gaussian Random Field Theory, and the correction for multiple comparisons family wise error correction (FWE).

Another approach to tackle the multiple comparison problem is to control the number of false-positive results (False Discovery Rate, FDR). Benjamini and Hochberg suggested to adjust the significance threshold in a way that the number of false positives in the thresholded map corresponds to the significance level (Benjamini and Hochberg 1995). Testing a dataset at a significance level of $p < 0.05$ FDR corrected means that no more than 5% of all above threshold voxels are false-positive results (Genovese et al. 2002).

Both methods (FDR and FWE) can be combined with an additional criterion, which takes the size of the activated clusters into account (Poline et al. 1997). Further approaches are permutation-based methods (Nichols and Holmes 2002) or simply restrict the search region to a (anatomically) well-defined and hypothesis-driven region of interest (ROI), which intrinsically reduces the number of comparisons performed. Further critical discussions of the statistical limitations in functional neuroimaging are presented elsewhere (Pettersson et al. 1999a, b).

12.6.4 Group Analysis

Up to this point, all analytical issues were based on single subject data, which were acquired in one run.

To analyze multiple runs or subjects, the GLM can be easily extended by adding new blocks to the model, with each block containing a run of one subject. This so-called fixed-effect analysis assumes hereby that β is constant across runs, session, and subject, and takes only intra-subject variability into account, which is attributed to “noise” in the fMRI data. Consequently, fixed-effect models only ask whether an effect can be observed within the examined cohort.

However, we are usually more interested in the question whether the observed effects are representative for the whole population from which the cohort was drawn. This can only be addressed if the inter-subject variability—the variability of the effect sizes between different subjects—is considered as well. This is done in the so-called random effect (RFX) analysis. The random effect analysis in its simplest form assumes that the effect of each subject follows a normal distribution. The variance is assumed to be constant. Additionally, it is assumed that population(s) is (are) normal distributed. Both assumptions are questionable and lead to overly conservative test results. This can be overcome by either using multilevel approaches (Beckmann et al. 2003; Friston et al. 2002; Woolrich et al. 2004a) or nonparametric techniques (Holmes et al. 1996; Nichols and Holmes 2002).

12.7 Advanced Methods

12.7.1 *Multivariate Analysis Methods*

The data analysis method that was described in the preceding sections was “model-” or “hypothesis-” driven: Assumptions about the shape of the hrf and the putative time course for an active voxel during the presentation of experimental conditions are formulated and fed into the GLM to compare the hypothesized with the observed values for each single voxel. However, this approach might not always be adequate. For example, the hrf might not be known for a patient population, or a paradigm might not be suited to be divided easily into different conditions (e.g., during a car driving simulation). In fact, in so-called resting brain fMRI experiments, which became popular in recent years (see below), there is hardly an explicit task at all (except for the instruction to lie still and stay awake).

Data-driven approach does not rely on such a priori models, and does not have to account for every possible effect. Multivariate methods do not work on the single voxel level, but take the whole dataset and try to find common structures in space, time, or both space and time. Since a model does not have to be specified in advance, these approaches might be able to discover the unexpected, but they have the disadvantage to be nonspecific. A significant finding could be image artifact as well as a real finding.

The most popular multivariate method at present is the independent component analysis (ICA). The key idea behind ICA is that a mixture of independent sources is “more statistically dependent” than the sources itself. For example, take the classical “cocktail party problem” (Fig. 12.6). Imagine two guests during a cocktail party (or in a noisy environment) who are speaking at the same time (processing on the behavioral and the systemic (activation) level. Two microphones record this “conversation”. Because the microphones are not exactly on the same position, both microphones record a slightly different mixture of the sources. ICA is capable to reconstruct the sources (= both speakers separately) from signals recorded by the microphones. If the conversations are denoted $S_1(t)$ and $S_2(t)$, and the recorded data $M_1(t)$ and $M_2(t)$, one could write

$$M_1(t) = a_{1,1}S_1(t) + a_{1,2}S_2(t), \quad (12.6)$$

$$M_2(t) = a_{2,1}S_1(t) + a_{2,2}S_2(t). \quad (12.7)$$

The coefficients a_{ij} form the mixing matrix \mathbf{A} . \mathbf{A} can be estimated with the assumptions that the $S_i(t)$ are non-normal, that the number of recorded tracks equals at least the number of independent components, and that the unmixing matrix consists of linear independent columns (Hyvarinen and Oja 2000). There are different approaches to solve this problem, but in the end, all are based on the idea that the

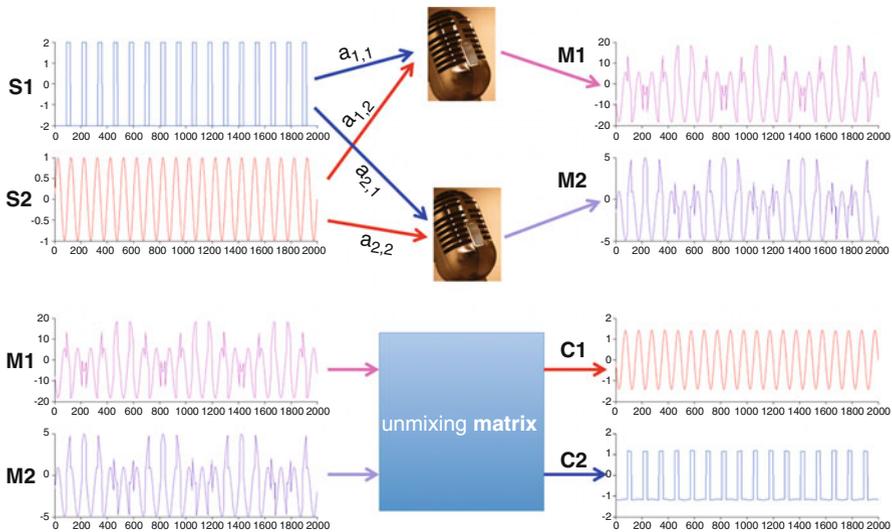


Fig. 12.6 *Cocktail Party Problem*. Two independent sources ($S1$, $S2$) are recorded by two independent microphones ($M1$, $M2$). They record different mixtures of both signals ($M1$, $M2$) depending on their position in space. Hereby it is assumed that the both signals mix linearly. Independent component analysis allows for calculating the so-called unmixing matrix, which is able to reconstruct two components ($C1$, $C2$) that represent the original sources ($S1$, $S2$). Please note that the numbering of the components does not (necessarily) correspond to the original numbering and that source and reconstructed independent component differ in amplitude and phase. This information is lost during the unmixing process

sources have to be less Gaussian than the mixed data. A nice introduction into the rationale of ICA can be found in a paper by Stone (2002).

ICAs are universally applicable, as long as the sources are linearly mixed and not perfectly Gaussian distributed. The result does not depend on the order of the entered data and reveals even useful results when the sources are correlated with each other. ICA can be used to detect and correct for artifacts in EEG data as well as to analyze fMRI data, or to find hidden factors in financial data. As a downside of this method, it has to be mentioned that ICA can only cope with linear mixtures of sources. Nonlinear interactions among modes are not detected by this method (see Friston 1998 for a critical discussion). It is also the job of the user to find the meaningful components and to identify the detected artifacts. In addition, ICA components have arbitrary amplitudes and even the sign is not defined, which means that a component might equally represent an activation or deactivation: In order to decide on the sign, one has to go back into the original data (see also Fig. 12.6).

Even though ICA can be used to analyze task fMRI data (McKeown et al. 1998), the most prominent use of ICA (in fMRI) is to analyze resting brain data. In resting brain experiments, volunteers are asked to relax and to keep their mind free of any thoughts (see also Chap. 16). By analyzing fMRI data acquired under such conditions, it has been shown that at least seven functionally and neuroanatomically

meaningful networks can be identified (Beckmann et al. 2005). These networks represent structures, which are anatomically connected, as can also be revealed using diffusion tensor imaging (Li et al. 2011; Skudlarski et al. 2008; Teipel et al. 2010, see below).

12.7.2 *Connectivity Analysis*

So far, we have primarily discussed how one could find certain brain regions. Unfortunately, focusing on functional segregation falls short of brain complexity. Every single brain region is embedded into a complex, interacting network. In order to understand brain function, functional pathology, or physiological changes due to interventions (which may also include exercise-related changes), it is also necessary to investigate the integration of brain regions into networks, as well as the interactions within and between networks. Currently, there are two distinct methods for characterizing the functional integration of brain areas within a network: functional and effective connectivity. Functional connectivity assesses functional integration by investigating the statistical dependencies between the activities of different brain regions. This is usually put into practice by calculating the cross-correlation coefficient between the time courses of preselected regions of interests. Even though functional connectivity is only descriptive, and it does not allow for distinguishing cause and effect, it is still very useful, for example, to classify groups. In this context, the correlation coefficient between well-defined regions can be used as an endophenotype distinguishing between populations (Craddock et al. 2009).

On the other hand, effective connectivity tries to directly access the influence between regions employing an explicit model of the coupling process. Effective connectivity is based on an anatomical model of regions that interact functionally and, in addition, on a model how this interaction takes place. In this sense, effective connectivity can be considered to be causative, whereas functional connectivity is more descriptive in nature. Yet, in both cases, the components of the network have to be specified. Currently, the most widely used approaches to access functional connectivity are structural equation modeling (SEM), dynamic causal modeling (DCM), and Granger causality. Within this troika, DCM is the most widely used method. Here a model is built, based on the neuronal level, and changes in neuronal activity are described by nonlinear differential equations. Applying a physiological model of the BOLD response (as the abovementioned balloon model) onto the neuronal model, the hidden neuronal states are transformed into a model for observable BOLD data. The connection strength between the a priori chosen regions is part of the model and can be estimated by comparing the predicted BOLD response with the experimental data. While the details cannot be discussed here, they are nicely covered in a recent review article by Karl Friston, in which the historical and conceptual, as well as the mathematical, background is discussed (Friston 2011).

Besides the above-mentioned approaches, there are a lot of alternative methods to analyze effective connectivity between brain regions, like coherence analysis,

generalized synchronization, or Patel's conditional dependence measure, and it is still an ongoing discussion which method might be best suited for analyzing brain connectivity. They differ in the sensitivity to detect network connections or directionality and therefore in their power to reveal causality within a given network (Smith et al. 2011). Interestingly, when comparing different modeling and analysis approaches, Smith and colleagues found that the most crucial parts in network modeling and effective connectivity analyses are the choice of functionally accurate ROIs. Connectivity analyses are extremely sensitive for inaccuracies in the initial step: network definition. Despite choosing the appropriate regions for the model itself, great care should be taken to identify these regions individually in each volunteer. The definition via a structural atlas seems not to be appropriate.

12.7.3 Network Analysis

An interesting question which is not addressed by the methods covered above is the *overall* network structure of the brain. How does information flow from one point to another? If information is transported between two regions, how many intermediate regions are needed? Do central relay structures exist, which are more important than others? Question of this kind can be investigated by a relatively new scientific field, called network science. Network sciences deal with properties of complex systems, and its applicability reaches from social sciences to physics, economy, and the description of ecosystems up to the evaluation of the properties of large neuronal cell populations as the human brain.

The mathematical theory behind network sciences is called graph theory. The mathematician Leonhard Euler founded graph theory when he tried to solve a popular puzzle in 1736. Königsberg, today Kaliningrad, was dissected into four parts by two branches of the river Pregel. Seven bridges connected the parts with each other. The question was whether a path exists that would allow make a walkabout, with each bridge being crossed only one time. Euler solved the problem by a graphical abstraction of the problem (Fig. 12.7). "Nodes" replaced the landmasses and "edges" replaced the bridges. By doing so, Euler was aware of the fact that the problem was independent of the actual distances or geographical shapes of Königsberg. Only the relative position of each part and its relation to the other part were of importance. He showed that the desired path does not exist and found a general solution that is applicable to an arbitrary number of nodes and edges.

Nodes and edges have a straightforward interpretation in the context of neuroscience: Nodes can be interpreted in terms of neurons, clusters of neurons, or brain regions, whereas edges are the connections between them. These connections can be directed or undirected and binary or weighted: as a consequence, directed/undirected and binary/weighted networks can be distinguished. Two different matrices, the connection matrix and adjacency matrix, are used to represent graphs mathematically. In these matrices, the columns and rows represent the nodes. For the connectivity matrix, the matrix entries represent the edges, and—in the case of the adjacency

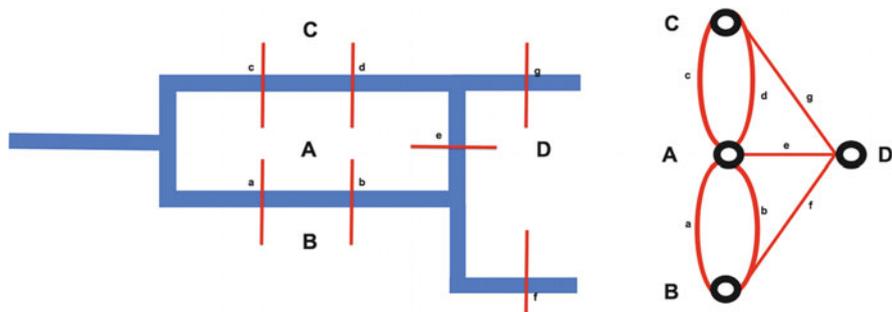


Fig. 12.7 *The Königsberg bridge problem.* Does a path exist that connects A–D via the bridges (red) a–g across the river Pregel (blue)? It is allowed to cross every bridge only one time and the end and starting point has to be the same. The geographical situation is schematically outlined on the left. Using the “graph representation” (drawn on the right), where the bridges are represented as a line, called “vertex”, and the landmasses as dark circles, called “nodes”, Leonard Euler was able to prove that such a path does not exist. He was even able to generalize the solution: A connecting path exists only if a graph contains exactly two nodes or no node with an uneven number of vertices. The drawings are based on Euler’s work

matrix—the minimal number of edges needed to connect two nodes. An ordered sequence of edges that links two nodes is called “path”, and the number of edges building a path is called “path length”. With these two measures, it is possible to calculate global and local properties like “degree”, which is the number of edges connected to a node or the average path length. Meanwhile, the adjacency matrix can be used to derive a third parameter, the “assortativity”, which measures the tendency of nodes to link to nodes with similar degree. In terms of neuronal networks, these measures have a direct interpretation. A (brain) region with a high “degree” is directly influenced by a large number of other regions. Path length and distance matrices characterize the communication structure within the network. Using the degree parameter, one can also identify those regions, which are most important for the overall functioning of the network, because these are the nodes (so-called hubs), which are most highly connected. But a high connectivity does not necessarily mean that its breakdown will critically disturb the functioning of the network as a whole. This will critically depend on the network’s overall architecture: If a network is built of subnetworks (“modes”), one can assume that the functioning of the network as a whole is not hampered, given that the critical hub only connects nodes within a mode. However, if a hub connects different modes, it might be crucial if it is taken out.

Networks can be classified with regarding to their construction process. Interestingly, this does also influence the properties of the network. There are two extremes: A purely random network, where each node is connected to other nodes by chance, and a regular lattice graph, which follows a predefined pattern (Fig. 12.8). Both types of networks have distinct features: Whereas random networks have small path lengths and a low level of clustering, the regular networks have a higher path length and a higher degree of clustering. It is not surprising that real-world networks

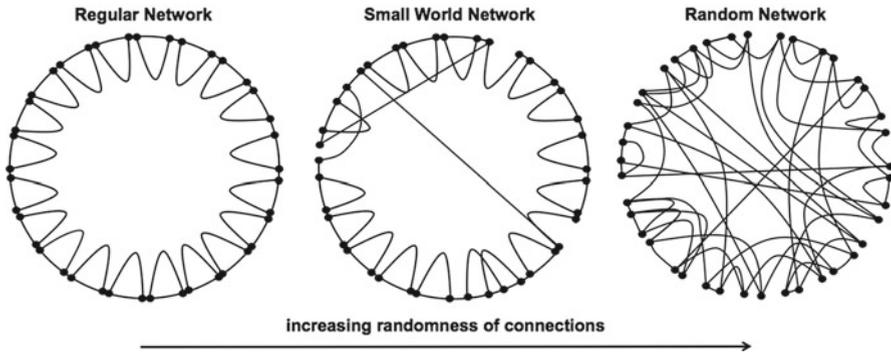


Fig. 12.8 *Network categories.* Three kinds of networks are distinguished: the regular network, the small world network, and the random network. By randomly rewiring a few connections, the regular network becomes a small-world network. If the number of randomly rewired connections increases, the network turns more and more a random network. The figure is adapted from Watts and Strogatz (1998)

are usually neither a pure random nor a regular lattice network; however, it is astonishing that most real-world networks seem to combine a high clustering coefficient with a short path length. Networks possessing this characteristic are called “small-world networks”. Duncan Watts and Stephen Strogatz discovered that by randomly rewiring a regular network, a high degree of clustering can be combined with a small path length (Watts and Strogatz 1998). This behavior already occurs at very small rewire probabilities. Practically, a small-world network can be constructed by taking a regular network and replacing a few edges connecting nearby nodes with edges connecting randomly chosen nodes (Fig. 12.8). As far as we know it today, these kinds of networks represent a large variety of real-world networks, ranging from the internet to social networks. Small-world network architectures can be found in the biochemical network of the living cell, in river networks, as well as in the brain. Within the context of this book, it is not possible to discuss small-world properties in detail. Instead, we like to refer to the literature. Besides the seminal paper of Watts and Strogatz (1998), a book written by Olaf Sporns can be recommended here (Sporns 2011). It is a good introductory book if one likes to understand the brain in the view of network theory. The mathematics behind it is nicely covered in numerous review articles, for example, Rubinov and Sporns (2010), Bullmore and Sporns (2009), or van den Heuvel and Hulshoff Pol (2010).

12.8 Practical Aspects of fMRI

When designing an fMRI experiment, one has to decide on the overall paradigm design, and it is difficult to give general recommendations, except that a paradigm should be as simple and short as possible. The easiest design type for sure is the block design. It is easy to implement and to analyze and has a high SNR. These aspects

are appealing and might guide the decision. However, the leading guideline should always be the scientific question behind the paradigm. If it is essential to distinguish between single events, a block design is simply not appropriate. It is also advantageous to collect behavioral data during the fMRI experiment. By doing so, it can be checked whether the experiment was successfully performed by a volunteer, and it helps to decide if any fMRI activation differences between groups actually rely on group membership, or need to be attributed to coexisting performance differences. Considering the costs of scan time (about 500\$/h), the paradigms should be tested beforehand. The positioning of the volunteer (supine), the restricted visual field, the handling of devices, age, and educational background of the volunteer (academic/nonacademic, children/adults/elderly etc.), and “click and feel” of the paradigm, should be evaluated during the testing phase. The optimum setup would be a mock scanner, which simulates the real scanner situation as close as possible (including scanner noise). The volunteers testing the final experiment should not be colleagues from the lab or students with a similar background, because they might anticipate the meaning of imprecise or misleading instructions, and might be familiar with experimental conditions, whereas a volunteer in general is not.

There are a plethora of software packages available for analyzing fMRI data (including preprocessing). Most of them are freely available and directly downloadable from the Internet. A good starting point is: <http://www.nitrc.org/>, where almost all freely available packages are linked. The most widely used packages are AFNI¹, BrainVISA², FSL³, SPM⁴, and the commercial package BrainVoyager⁵. Most vendors of MR scanners also have basic software packages available, but they usually provide only very basic preprocessing and analysis methods, which are rather designed for individual analyses than for group studies.

An important point that should be kept in mind when interpreting imaging results is that the large number of independent comparisons performed when analyzing data voxel-wise will produce false-positive results, no matter which correction method is used. Even if the location of activation fits perfectly into an a priori hypothesis, it does not automatically exclude the possibility that this finding might be a false-positive result. This was nicely illustrated by Craig Bennett and colleagues showing a brain activation of “dead salmon” (Bennett et al. 2009). The main lesson to learn from this paper is that one should not base a larger project on a single fMRI experiment, and consider a test-retest design, to validate the own data. Of course, this might not always be possible, but at least a split-half analysis may be feasible, especially when a larger project is planned based on a single fMRI experiment. A more general discussion on the reliability of fMRI results might also be found in a recent paper by Bennett and Miller (2010).

¹ <http://afni.nimh.nih.gov/afni>.

² <http://brainvisa.info/>.

³ <http://www.fmrib.ox.ac.uk/fsl/>.

⁴ <http://www.fil.ion.ucl.ac.uk/spm/>.

⁵ <http://www.brainvoyager.com/>.

12.9 Structural Analysis of the Brain

MRI allows for the acquisition of high-resolution structural images with an excellent image contrast between white matter, gray matter, and cerebrospinal fluid. This enables researchers to investigate structural differences (volume or shape of anatomical structures) between populations of interest. Until the development of automatic methods, morphometric analyses were performed by delineating anatomical structures manually, which is of course an extremely time-consuming and personnel-intensive business. As a consequence, those analyses were limited to small group sizes and to well-defined anatomical structures like the hippocampus, the ventricles, or the corpus callosum. An exploratory data analysis on large datasets is generally not feasible. This changed when automatic methods like voxel-based morphometry became available. Here, one can distinguish between voxel-based and surface-based methods (Ashburner et al. 2003). The development of diffusion imaging methods has opened additional venues for the analysis of structural differences in the white matter, which extend beyond classic volumetric approaches, by either analyzing the fractional anisotropy or the structural connectivity by reconstructing white matter tracts. The following sections introduce the key ideas of the most widely used morphometric methods.

12.9.1 Manual Morphometry

Manual morphometry is based on the manual delineation of anatomical structures on MR images and can be considered the classical region of interest (ROI) approach. In general, it can be applied to every brain structure that can be unambiguously identified on structural images. The key point here is the definition and standardization of landmarks that establish the border to neighboring structures. The easiest measure here is the volume and the surface of a structure. Another popular measure is the gyrification index, which compares the inner contour length following the sulci/gyri with the outer contour length enfolding the outer surface of the gyri.

Several software packages are available which can be used for this purpose (i.e., free packages: MRICron⁶, Mango⁷, MeVisLab⁸, fslview⁹, or, as a commercial package, ANALYZE¹⁰). Yet, it should be pointed out that manual morphometry is not only extremely time-consuming, but also personnel-intensive, because more than one person should perform the measurements. The researchers should be rigorously trained to ensure a high (inter-)rater reliability (e.g., Inter-rater Correlation Coefficient

⁶ <http://www.cabiatl.com/mricro/mricron/index.html>.

⁷ <http://ric.uthscsa.edu/mango/>.

⁸ <http://www.mevislab.de/>.

⁹ <http://www.fmrib.ox.ac.uk/fsl/fslview/index.html>.

¹⁰ <http://www.mayo.edu/bir/Software/Analyze/Analyze.html>.

ICC>0.85). The operators who conduct the analyses must be blinded for subject characteristics (and investigated hemisphere, if possible). The reliability should be reestablished during ongoing studies, and all definitions should be checked by at least one other person knowing the procedure. Usually, editors and reviewers want to see the actual quality control numbers if they are familiar with manual morphometry.

Despite the fact that manual morphometry demands a lot of resources, it also has a couple of advantages compared with automated measures. For example, it acts in native space of the data, which means that the datasets do not have to be spatially transformed or preprocessed, as it is the case for the automated methods (see below). It delivers unambiguous measures like volumes or surface areas and has, by definition, a high anatomical validity, which still has to be proven for many automatic methods. To conclude, if the resources are available and a clear a priori hypothesis can be defined, manual morphometry should be the method of choice to access structural differences in the brain.

12.9.2 Automated Methods

The automated methods have the advantage that almost no user interaction is necessary to analyze morphological differences between or within populations. Because no anatomical a priori hypothesis is needed, these methods can be used to conduct exploratory analyses of huge datasets. Meanwhile, unlike manual morphometry, automated morphometric methods do not work on data in native space. They have to be preprocessed before the analysis can be conducted. The preprocessing stream for all methods is more or less the same: skull stripping, intensity normalization to remove spatial intensity inhomogeneities, segmentation into gray matter (GM), white matter (WM) and cerebrospinal fluid (CSF), and normalization into a common space. Applying filters like anisotropic diffusion filters or taking into account neighborhood relations might additionally improve the segmentation quality (Cuadra et al. 2005). The validity of the results crucially depends on the segmentation quality. This holds true for all automated methods.

12.9.2.1 Voxel-Based Morphometry

Voxel-based morphometry (VBM) was the first automated method that was available to the neuroscience community. Though it was introduced 15 years ago by Wright and colleagues (1995), it took about 5 years before it became a popular methodology for assessing structural brain differences. Its popularity increase might be attributed to the publication of three papers: Two of them were covered extensively by the popular press, showing that even the adult human brain is capable of structural adaption to external challenges. The first one indicated that the hippocampal volume of taxi drivers increases depending on the amount of time spent as a taxi driver (Maguire et al. 2000), and the other one proved that even a short intervention of

jiggling training over 3 months leads to a transient GM volume increase in regions relevant for visuo-motor processing (Draganski et al. 2004). The third paper published by Ashburner and Friston (2000) extensively explained the method and made clear that the whole procedure could be performed within the framework of SPM (one of the most widely used fMRI preprocessing and analysis packages). Further programming was neither needed, nor special MRI sequences were used. The method works well with conventional isotropic T1-weighted datasets, which are usually acquired along with fMRI experiments.

When segmenting brain images into its three classes (GM, WM, CSF), one assigns each voxel a probability of belonging to one of these classes. Voxel-based morphometry is based on the idea that these probability maps of the segmented brain images contain local volume information. As outlined by John Ashburner and Karl Friston, these probability maps can be used to make voxel-wise comparisons of local volume differences across groups, given that gross anatomical differences between the brains are compensated (Ashburner and Friston 2000). The latter can be realized by the normalization process, which was already outlined in Sect. 12.5.4 for functional imaging data. Good et al. optimized this method by recursively constructing a study-specific template and by adding an additional preprocessing step, called “modulation” (Good et al. 2001). Modulation compensates for local volume changes by scaling the image intensities by the amount of contraction or dilatation applied when normalizing the dataset. By doing so, the total amount of gray matter remains the same as in the original image. Figure 12.9 outlines the processing stream of VBM and “optimized VBM”. The most recent improvement in VBM was to increase the amount of deformations allowed during the normalization process, which improves the realignment quality of small structures (Ashburner 2007). The newly developed algorithm is called Diffeomorphic Anatomical Registration using Exponentiated Lie algebra (DARTEL). Figure 12.10 compares the results of both methods (VBM with and without DARTEL).

12.9.2.2 Deformation-Based Morphometry

A further development of the VBM method is the deformation-based morphometry (DBM). It differs from VBM in the following points: the spatial normalization into a common space is performed using nonlinear deformations, and DBM does not compare the obtained normalized probability map for GM, WM, or CSF, but the transformation fields which are used to normalize the data (Ashburner et al. 1998). Because these deformations generally correct for global size and shape differences, DBM acts on a different spatial scale as compared to VBM. Therefore, it can be considered as complementary to VBM. Because deformation fields are multidimensional measures, a multivariate approach has to be applied to analyze the data statistically. This can be avoided by extracting a scalar measure from the deformation field. A common way is to calculate the determinant of the Jacobian matrix. The Jacobian matrix is built by the partial derivatives of a vector field. Its determinant is a measure for the magnitude of local deformations. This approach is also referred to in the literature as tensor-based morphometry (TBM).

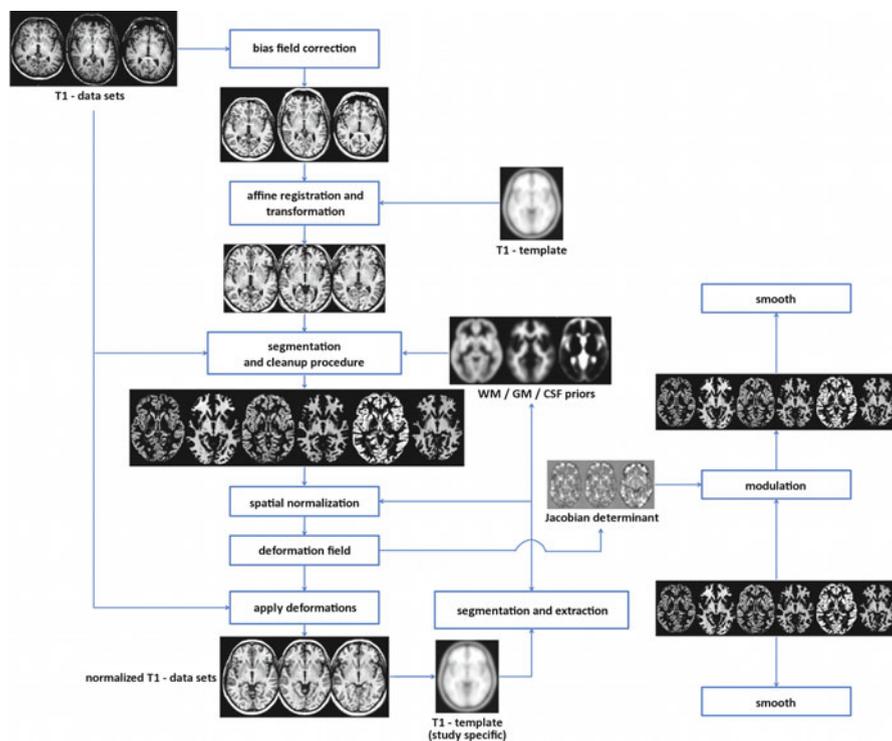


Fig. 12.9 *Processing scheme optimized VBM.* The key idea behind optimized VBM is to improve the quality of the segmentation and the normalization process. This is achieved by an iteratively constructed study-specific T1-template, and probability maps for GM and WM segmentations (GM-/WM-priors). An additional improvement of spatial normalization results is achieved by using the GM/WM segmentations for normalization procedure, instead using solely a T1-weighted dataset itself. After this initial normalization and segmentation, a study-specific template and segmentation class priors can be constructed. They are then used for the final normalization and segmentation. The “modulation” step (=voxel-wise multiplication with the Jacobian determinate) keeps the total amount of GM/WM constant. The bias-field correction corrects for large-scale intensity variation caused by inhomogeneous RF-deposition, which especially is seen at higher field strengths. The cleanup procedure(s) consist(s) of spatial filters. Both procedures (bias-field correction and cleanup) additionally improve the segmentation quality. The figure is adapted from Good et al. (2001)

12.9.2.3 Surface-Based Morphometry

Voxel-based or deformation-based methods measure morphological changes indirectly by analyzing either probability maps of a segmentation class or deformation fields. A more direct approach is to access geometric measures like cortical thickness, curvature, or sulcal depth. For this purpose, it is necessary to reconstruct the cortical sheet. It is defined by two boundaries: the inner boundary between the WM and GM, and the outer boundary between GM and CSF. The outer boundary corresponds anatomically to the pial surface, which can be reconstructed after successful

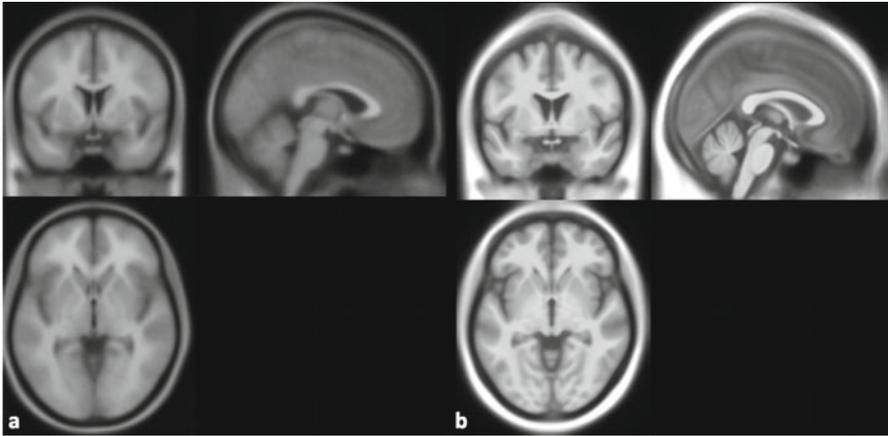


Fig. 12.10 Comparison between two normalization procedures. Resulting mean datasets after averaging the normalized T1-datasets of 20 healthy volunteers following either the “optimized-VBM” protocol (a) or the high-dimensional DARTEL normalization procedure (b). Please note the fine anatomical details preserved in (b)

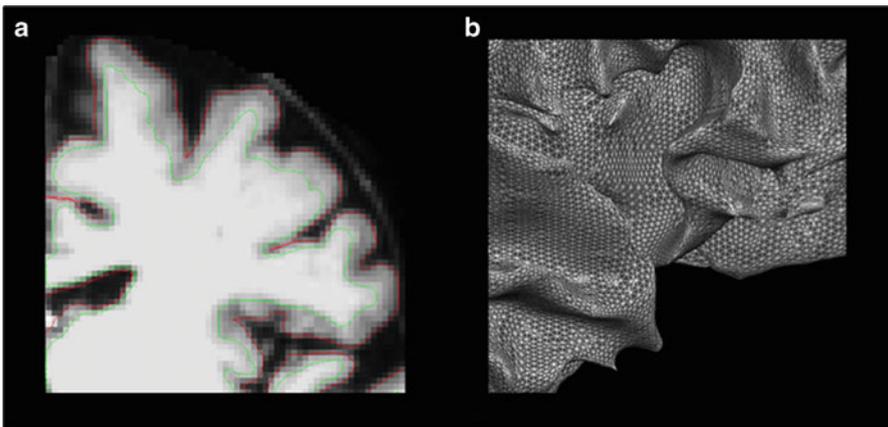


Fig. 12.11 Mesh reconstruction. Determination of the WM–GM (red) and GM–CSF (green) boundary by FreeSurfer (right). On the left, the triangulation of the WM surface is shown

skull stripping using the “Brain Extraction Tool” (BET, Smith 2002) or more advanced methods that combine watershed algorithms and deformable surface models (Segonne et al. 2004). The white matter is segmented based on its intensity. When both boundaries are determined, the surfaces can be parameterized using triangulation. Doing so, a dense mesh is constructed around the cortical sheet (Dale et al. 1999; see also Fig. 12.11). Using this mesh, measures like cortical thickness, curvature, or gyrification index can be directly calculated.

Besides the availability of direct measures for local and global geometric properties of the cortex, the reconstruction of the cortical sheet has an additional advantage. It allows for inter-subject registration using a surface-based coordinate system (Fischl et al. 1999a), which is more accurate compared with the voxel-based methods described above.

12.9.2.4 Practical Aspects

The preceding section gave an overview on the most common analysis approaches and is far away from being complete. But which method should be used? The gold standard is still manual morphometry. The advantage is its anatomical validity, and that it does not request any kind of preprocessing. However, it is extremely time-consuming and should be based on a priori hypotheses for anatomically well-defined structures. Usually, only small cohorts can be analyzed using this method. The automatic methods have the advantage that they are fast and need less human resources. They work on the whole brain or, at least, on the whole cortex. As a consequence, it is not necessary to define a priori hypotheses and methods can be applied to large cohorts, even in an exploratory manner. Cortex-based analyses provide anatomically meaningful measures like cortical thickness or cortex curvature. Group averaging can be performed on the cortical surface minimizing the sulcal variability within and across groups. Therefore, any effect that is found in cortical thickness maps (or curvature, etc.) can be clearly attributed to differences in this measure. Voxel-based and deformation-based morphometry are less specific in this respect. They use indirect measures, like probability maps or deformation fields, to find structural differences. Even though most authors of morphometry papers refer to GM/WM density, these methods do not assess GM/WM density or volume, but assess the probability for a certain voxel to be GM/WM: This is an ambiguous information, because this probability does not only depend on GM/WM content but also on the average shape of this structure across groups. However, VBM has the practical advantage that it can be used within the software used for fMRI analyses, and it is much faster as compared to the cortex-based methods. Meanwhile, VBM has the disadvantage that it is not optimal for longitudinal studies. As a recent study showed, drift effects in scanner hardware and inter-scanner variability lead to global and regional effects mimicking real brain volume changes (Takao et al. 2011). Surface-based and deformation-based methods seem to be less vulnerable to such effects. Both methods can also be applied to longitudinal case studies whereas VBM generally needs large cohorts.

However, it should be also noted at this point that if one is specially interested in a certain structure, optimized methods might exist that outperform the approaches outlined above. Most of them are semi-automated, combining manual delineation of a structure with automated analysis protocols. For example, there are specialized methods for analyzing the regional hippocampal thickness (Zeineh et al. 2001) or shape (Tepes et al. 2008).

Software for morphometric analysis is freely available in most of the cases. Voxel-based morphometry can be performed with any software package that

provides segmentation and normalization, like SPM, FSL, AFNI, or BrainVoyager. Surface-based analyses can be performed, for example, by FreeSurfer¹¹, CARET¹², mrVista/itkGray¹³, SUMA¹⁴, BrainVoyager, or BrainVISA/Anatomist¹⁵. The latter also contains toolboxes for sulcal recognition and morphometry, including gyrification index, sulcal length, and depth.

12.9.3 Analysis of Diffusion-Weighted Images

Diffusion-weighted imaging (DWI) permits the investigation of microstructure and integrity of the white matter and its anatomical connectivity. As pointed out in Chap. 11, it is based on the diffusion (= random motion or Brownian motion) of water molecules. Since the motion of water in axons is restricted perpendicular to the main direction of an axon, the main direction of axon bundles can be inferred by applying multiple diffusion gradients in different directions and strengths. Unfortunately, running diffusion-weighted MR sequences do not result in a single image volume that already contains parameters for diffusion direction and strength on a voxel-by-voxel basis. Instead, one obtains a series of image volumes. The amount of diffusion weighting in one image volume depends on the strength of the diffusion gradient applied when the volume was acquired and its directional sensitivity in the gradient direction. In the following sections, it will be explained how diffusion-related quantities are calculated and analyzed. As in the preceding sections, the following discussion has to be restricted to the main ideas. A nice introduction into the mathematics in diffusion tensor imaging can be found in three papers written by Kingsley (2006a, b, c).

12.9.3.1 Calculation of Diffusion Properties

The diffusion of water molecules in a homogenous volume does not prefer any spatial direction. The motion is isotropic. If the water is confined in a small tube, this motion is restricted perpendicular to the long axis of this tube, whereas it is unrestricted parallel to the long axis. The diffusion motion can be quantified using the diffusion coefficient as introduced by Einstein. The mean displacement $\langle X_i^2 \rangle$ of molecules in a liquid three-dimensional space over time t can be expressed as

$$\langle X_i^2 \rangle = 6 * D * t, \quad (12.8)$$

where D denotes the diffusion coefficient, and X_i the displacement vector of the i th molecule in the liquid. In the isotropic case, no preferred direction exists. If all displacement vectors X_i would be drawn after the time t , they would approximate a

¹¹ <http://surfer.nmr.mgh.harvard.edu/>.

¹² http://brainvis.wustl.edu/wiki/index.php/Main_Page.

¹³ <http://white.stanford.edu/software/>.

¹⁴ <http://afni.nimh.nih.gov/afni/suma>.

¹⁵ <http://brainvisa.info/>.

sphere whose radius depends on D and t . However, if the diffusion was restricted, i.e., by a tube (like an axon), the displacement vectors would shape an ellipsoid with its long axis parallel to the long axis of the tube, because the diffusion is restricted perpendicular to this axis. While it is not sufficient to use a single value to characterize the diffusion, however, it would be sufficient to measure the diffusion coefficient along the three main axes of the ellipsoid. In the more general case, a diffusion tensor \mathbf{D} is needed to describe the diffusion:

$$\mathbf{D} = \begin{pmatrix} d_{xx} & d_{xy} & d_{xz} \\ d_{yx} & d_{yy} & d_{yz} \\ d_{zx} & d_{zy} & d_{zz} \end{pmatrix} \quad (12.9)$$

The indices refer to the axes of any user-defined coordinate system ($xx=x$ -axis, xy =diagonal between x - and y -axis, etc.). Because \mathbf{D} is symmetric ($d_{xz}=d_{zx}$...), it is sufficient to determine six elements to construct the complete tensor. This can be afforded by applying diffusion gradients in six different directions. Having done this, the diffusion is specified in the frame of reference defined by the gradient directions (usually the scanner reference frame). However, it can be mathematically transformed into the optimal frame of reference, where the coordinate axes are oriented parallel to the main- or principal axis of the diffusion ellipsoid.

$$\mathbf{D} = \begin{pmatrix} d_{xx} & d_{xy} & d_{xz} \\ d_{yx} & d_{yy} & d_{yz} \\ d_{zx} & d_{zy} & d_{zz} \end{pmatrix} \Rightarrow \text{diagonalize} \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix}. \quad (12.10)$$

The values λ_1 , λ_2 , and λ_3 refer to the diffusibility along the principal diffusion directions A_1 , A_2 , and A_3 . They are determined when \mathbf{D} is diagonalized. Practically, this means that one can determine direction and amount of the diffusion when the diffusion is measured in (at least) six directions.

Everything stated so far does not completely account for the situation in biological tissue. Additional components exist which influence the molecular motion, like capillary flow or plasma flow in axons. Because these effects cannot be separated from the Brownian motion, the measured diffusion coefficients are considered as “apparent diffusion coefficient”, and called ADC accordingly.

12.9.3.2 Quantitative Diffusion Maps

The diagonalized diffusion tensor allows for calculating two measures, the fractional anisotropy (FA) and the mean diffusibility (MD). They are defined as follows:

$$\text{FA} = \sqrt{\frac{3 \sum_{i=1}^3 (\lambda_i - \bar{\lambda})^2}{2 \sum_{i=1}^3 \lambda_i}} \quad (12.11)$$

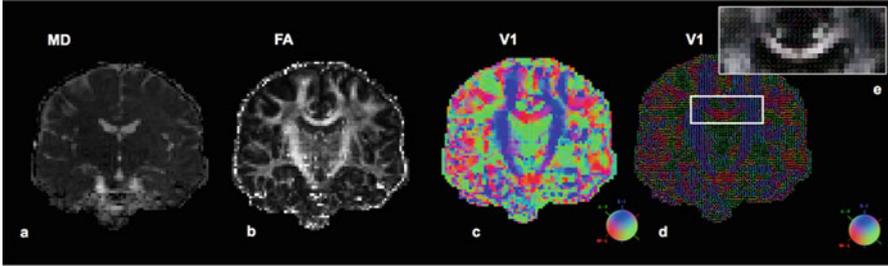


Fig. 12.12 *Reconstruction of different diffusion measures.* Mean diffusivity (a), fractional anisotropy (b), and the first eigenvector (V1) of the diffusion tensor are shown (c/d). The orientation of the eigenvector can be either visualized by using a certain colors for certain direction (red: medial—lateral, blue: superior—inferior, green: anterior—posterior) as outlined in (c), or by representing the directions by “lines” as it is commonly done for “vector fields”. “Lines” are often used instead of “vectors”, because only the orientation in space is meaningful (and not the direction). (d) Combines both approaches. The “vector field” representation can be superimposed onto structural data or, as shown in (e), onto any diffusion measure. Here, the first eigenvector is projected onto the FA map. Please note the high anisotropy of the capsula interna and the corpus callosum

and

$$\text{MD} = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3}. \quad (12.12)$$

FA can take values between 0 (isotropy) and 1 (maximal anisotropy). It measures how strongly directional diffusion is in a given voxel, and is thought to provide a good marker for WM integrity. In contrast, MD is simply a measure of the amount of diffusion within a voxel irrespective of direction (Fig. 12.12).

A straightforward approach for the within- or across-group analysis of FA and MD maps seems to be the application of those procedures that we already have described for VBM: normalization, smoothing, and univariate voxel-wise analysis. Indeed, this procedure has already led to valuable results (for example, Buchel et al. 2004), but it is problematic because the alignment of corresponding structures in the white matter (=fiber bundles) is very difficult across groups. Even small spatial shifts across subjects in a group can reduce the mean FA, which are then misinterpreted as real FA. Smoothing could help in this situation, but it is difficult to determine the amount of smoothing needed. One way to overcome this problem is a ROI-based approach, or to analyze the FA or MD values along reconstructed fiber tracts (see below). However, both methods do not allow for a whole brain analysis. An alternative method that was developed by Steven Smith and colleagues (Smith et al. 2006) is called tract-based spatial statistics (TBSS). The main idea is to perform voxel-wise cross-subject statistics on a WM skeleton, which represents the centers of all fiber tracts. These skeletons are constructed on the basis of the FA maps. After construction of these skeletons, the mean FA values can be projected onto them and compared across groups (Fig. 12.13).

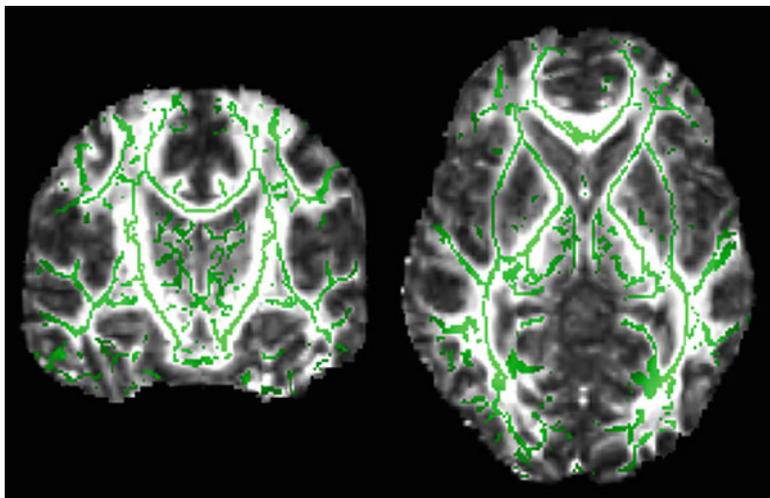


Fig. 12.13 *Skeleton*. Reconstruction of the “skeleton” for the main white matter tracts (*green*). As it can be depicted nicely from the images, the skeleton projects always onto the middle of the large fiber tracts. The FSL course test datasets were used to generate these images. (http://www.fmrib.ox.ac.uk/fslcourse/fsl_course_data2.tar.gz)

12.9.3.3 Analysis of Diffusion Tensors

So far, only scalar measures were used, ignoring the directional information that the diffusion tensor contains. This information can be used to reveal the large fiber bundles in the brain, assuming that the direction of the maximum diffusivity represents the major fiber orientation within a voxel. By plotting the principal diffusion direction (which can be indicated by a vector or using a color coding scheme), the gross structure of the white matter can be visualized on a voxel-by-voxel basis (Fig. 12.12).

Whereas these maps give a broad overview about the bundle orientation within the white matter, it is difficult to extract concrete information on single fiber bundles. This can be achieved by starting at a seed point and following the main direction voxel by voxel until it reaches gray matter. Mathematically this problem can be expressed by a differential equation:

$$\frac{d\vec{r}(s)}{ds} = \vec{v}(\vec{r}(s)), \quad (12.13)$$

where \vec{r} denotes a position along a curve and \vec{v} the main principal direction. Usually, additional boundary conditions are employed to avoid wrong assignments, like FA thresholds to ensure that the “fibers” only propagate in regions with meaningful anisotropy, or a maximum curvature to avoid kinks and extremely bended paths. The success of this “deterministic fiber tracking” approach to extract the main fiber bundles is nicely shown by Catani and Thiebaut de Schotten (2008).

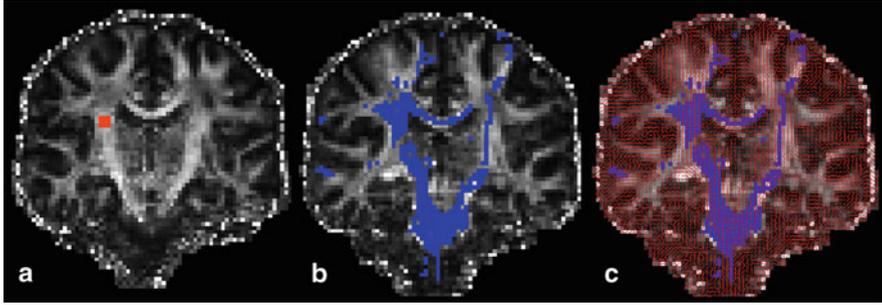


Fig. 12.14 *Tensor-tracking*. Using seed voxel(s) (a), the major connections (blue) between the seed voxel(s) and the rest of the WM voxels can be reconstructed with probabilistic fiber tracking approaches (b, c). In comparison to simple approaches that follow only the main directions within the tensor field (not shown), the probabilistic tracking approach also finds more complex pathways like the fibers that cross the corpus callosum and/or run through the contralateral capsula interna. The relation between the tensor field and the tracks is shown in (c)

An alternative approach to derive white matter tracts is probabilistic tracking, which assigns each diffusion direction in a given voxel with a probability. The probability is based on the directional information of the tensor and the measurement error of the tensor. The error term describes the uncertainty of the calculated main diffusion direction. Combining both parameters, one can construct all possible paths starting from a given seed region. It results in a probabilistic map that assigns every voxel (in the brain) a probability of being reached via a fiber bundle running through the seed region (Behrens et al. 2003b; Parker and Alexander 2003; Parker et al. 2003; see also Fig. 12.14). One of the advantages of the probabilistic tracking method is that it also allows to access regions with low FA values, and that it does not rely on artificial thresholds like the FA threshold or maximum-curvature criterion, as it can be found for deterministic tractography. By comparing all paths through a voxel with all possible paths, also quantification is possible. Moreover, because no FA threshold has to be applied, the method can be used to track the white matter paths right into the gray matter: This has been impressively shown by Tim Behrens, Heidi Johanson-Berg, and colleagues, who were able to distinguish thalamic subnuclei by analyzing their cortical connectivity (Behrens et al. 2003a).

12.9.3.4 Crossing Fibers

One issue that we have not discussed so far is the problem of ambiguity of the diffusion tensor in regions with crossing or kissing fiber bundles. At the commonly used spatial resolution ($2 \times 2 \times 2 \text{ mm}^3$), DTI does not resolve a single fiber but average across a vast amount of axons. The calculation based on the measurement of six diffusion directions does not account for crossing or kissing fibers, which can lead to wrong path assignments or artificially low FA values. This issue can be resolved by measuring a higher number of diffusion directions. This allows for calculating higher-order tensor terms that can account for additional preferred diffusion directions.

The more diffusion directions are acquired, the better crossing fiber regions can be evaluated. However, it increases also the measurement time and, hence, the probability of motion artifacts (Alexander 2005).

12.9.3.5 Practical Aspects

A couple of practical aspects have to be considered when DTI is performed. The strength of the diffusion gradients should be chosen between 1,000 and 1,500 s/mm². To compensate for motion, it is of advantage to acquire additional volumes without diffusion weighting. It also has to be decided how many directions have to be measured. The optimum number of diffusion directions is still under debate. To calculate FA values, 6–12 directions or more give stable results. For tractography, 30–60 diffusion directions seem to be appropriate, but this depends on the scientific question: The major fiber bundles can be already tracked using 12–30 diffusion directions, but for probabilistic tractography frequently a higher number of directions are used (60+). Independent of the number of directions, the direction should be evenly distributed on a sphere. The optimal parameters are also hardware dependent and should be discussed with the local physicist group.

The next obstacle when doing DTI is to extract the diffusion direction. The diffusion directions are needed in image space. The directions cannot be directly gained from the sequence protocol, because the image space depends on the patient position, whereas the gradient directions are defined using the scanner coordinate system as frame of reference. However, the information can be directly extracted from the DICOM¹⁶ header of the DTI volumes, i.e., using MRICron¹⁷. Most software packages that are capable of extracting the diffusion direction (and gradient amplitude) on the basis of DICOM data are also able to convert the DTI data into an appropriate data format, like Nifti¹⁸. The preprocessing steps (motion correction and so-called eddy-current compensation) can be performed within the same framework as the calculation of the ADC and FA maps. The most widely used free software packages for DTI analyses are FSL, DTI-studio, and BrainVISA, or BrainVoyager as the commercial counterpart. A good starting point to look for tools and free software packages is <http://www.nitrc.org>.

References

- Aguirre GK, Zarahn E, D'Esposito M (1998) The variability of human, BOLD hemodynamic responses. *Neuroimage* 8:360–369
- Aguirre GK, Detre JA, Zarahn E, Alsop DC (2002) Experimental design and the relative sensitivity of BOLD and perfusion fMRI. *Neuroimage* 15:488–500

¹⁶ DICOM is a standard data format for medical images and available on every MR-scanner.

¹⁷ <http://www.cabiatl.com/mricro/mricron/index.html>.

¹⁸ <http://nifti.nimh.nih.gov/>.

- Akgoren N, Dalgaard P, Lauritzen M (1996) Cerebral blood flow increases evoked by electrical stimulation of rat cerebellar cortex: relation to excitatory synaptic activity and nitric oxide synthesis. *Brain Res* 710:204–214
- Alexander DC (2005) Multiple-fiber reconstruction algorithms for diffusion MRI. *Ann N Y Acad Sci* 1064:113–133
- Amaro E Jr, Barker GJ (2006) Study design in fMRI: basic principles. *Brain Cogn* 60:220–232
- Andersson JL, Hutton C, Ashburner J, Turner R, Friston K (2001) Modeling geometric deformations in EPI time series. *Neuroimage* 13:903–919
- Ashburner J (2007) A fast diffeomorphic image registration algorithm. *Neuroimage* 38:95–113
- Ashburner J, Friston KJ (2000) Voxel-based morphometry—the methods. *Neuroimage* 11:805–821
- Ashburner J, Friston K (2007a) Non-linear registration. In: Friston KJ, Ashburner J, Kiebel S, Nichols TE, Penny W (eds) *Statistical parametric mapping: the analysis of functional brain images*. Academic, London, pp 63–80
- Ashburner J, Friston KJ (2007b) Rigid body registration. In: Friston KJ, Ashburner J, Kiebel S, Nichols TE, Penny W (eds) *Statistical parametric mapping: the analysis of functional brain images*. Academic Press, London, pp 49–60
- Ashburner J, Hutton C, Frackowiak R, Johnsrude I, Price C, Friston K (1998) Identifying global anatomical differences: deformation-based morphometry. *Hum Brain Mapp* 6:348–357
- Ashburner J, Csernansky JG, Davatzikos C, Fox NC, Frisoni GB, Thompson PM (2003) Computer-assisted imaging to assess brain structure in healthy and diseased brains. *Lancet Neurol* 2:79–88
- Attwell D, Iadecola C (2002) The neural basis of functional brain imaging signals. *Trends Neurosci* 25:621–625
- Attwell D, Buchan AM, Chrapak S, Lauritzen M, Macvicar BA, Newman EA (2010) Glial and neuronal control of brain blood flow. *Nature* 468:232–243
- Bandettini PA, Wong EC, Hinks RS, Tikofsky RS, Hyde JS (1992) Time course EPI of human brain function during task activation. *Magn Reson Med* 25:390–397
- Bandettini PA, Wong EC, DeYoe EA, Binder JR, Rao SM, Birzer D (1993) The functional dynamics of blood oxygen level dependent contrast in the motor cortex. *Proc Int Soc Magn Reson Med*, New York, p 1382
- Beckmann CF, Jenkinson M, Smith SM (2003) General multilevel linear modeling for group analysis in fMRI. *Neuroimage* 20:1052–1063
- Beckmann CF, DeLuca M, Devlin JT, Smith SM (2005) Investigations into resting-state connectivity using independent component analysis. *Phil Trans R Soc Lond B Biol Sci* 360:1001–1013
- Behrens TE, Johansen-Berg H, Woolrich MW, Smith SM, Wheeler-Kingshott CA, Boulby PA, Barker GJ, Sillery EL, Sheehan K, Ciccarelli O, Thompson AJ, Brady JM, Matthews PM (2003a) Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nat Neurosci* 6:750–757
- Behrens TE, Woolrich MW, Jenkinson M, Johansen-Berg H, Nunes RG, Clare S, Matthews PM, Brady JM, Smith SM (2003b) Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn Reson Med* 50:1077–1088
- Behzadi Y, Liu TT (2006) Caffeine reduces the initial dip in the visual BOLD response at 3 T. *Neuroimage* 32:9–15
- Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Series B (Methodological)* 57:289–300
- Bennett CM, Miller MB (2010) How reliable are the results from functional magnetic resonance imaging? *Ann N Y Acad Sci* 1191:133–155
- Bennett CM, Miller MB, Wolford GL (2009) Neural correlates of interspecies perspective taking in the post-mortem Atlantic Salmon: an argument for multiple comparisons correction. *NeuroImage* 47(Suppl 1):S125
- Blamire AM, Ogawa S, Ugurbil K, Rothman D, McCarthy G, Ellermann JM, Hyder F, Rattner Z, Shulman RG (1992) Dynamic mapping of the human visual cortex by high-speed magnetic resonance imaging. *Proc Natl Acad Sci USA* 89:11069–11073

- Boynton GM, Engel SA, Glover GH, Heeger DJ (1996) Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16:4207–4221
- Brett M, Johnsrude IS, Owen AM (2002) The problem of functional localization in the human brain. *Nat Rev Neurosci* 3:243–249
- Brewer JB, Zhao Z, Desmond JE, Glover GH, Gabrieli JD (1998) Making memories: brain activity that predicts how well visual experience will be remembered. *Science* 281:1185–1187
- Buchel C, Raedler T, Sommer M, Sach M, Weiller C, Koch MA (2004) White matter asymmetry in the human brain: a diffusion tensor MRI study. *Cereb Cortex* 14:945–951
- Buckner RL, Bandettini PA, O'Craven KM, Savoy RL, Petersen SE, Raichle ME, Rosen BR (1996) Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging. *Proc Natl Acad Sci USA* 93:14878–14883
- Bullmore E, Sporns O (2009) Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci* 10:186–198
- Buxton RB (2001) The elusive initial dip. *Neuroimage* 13:953–958
- Buxton RB, Frank LR (1997) A model for the coupling between cerebral blood flow and oxygen metabolism during neural stimulation. *J Cereb Blood Flow Metab* 17:64–72
- Buxton RB, Wong EC, Frank LR (1998) Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. *Magn Reson Med* 39:855–864
- Catani M, Thiebaut de Schotten M (2008) A diffusion tensor imaging tractography atlas for virtual in vivo dissections. *Cortex* 44:1105–1132, A journal devoted to the study of the nervous system and behavior
- Chein JM, Schneider W (2003) Designing Efficient fMRI Experiments. In: Grafman J, Robertson IH (eds) *Handbook of neuropsychology*. Elsevier, Amsterdam, pp 299–326
- Craddock RC, Holtzheimer PE 3rd, Hu XP, Mayberg HS (2009) Disease state prediction from resting state functional connectivity. *Magn Reson Med* 62:1619–1628
- Cuadra MB, Cammoun L, Butz T, Cuisenaire O, Thiran JP (2005) Comparison and validation of tissue modelization and statistical classification methods in T1-weighted MR brain images. *IEEE Trans Med Imaging* 24:1548–1565
- Dale AM, Buckner RL (1997) Selective averaging of rapidly presented individual trials using fMRI. *Hum Brain Mapp* 5:329–340
- Dale AM, Fischl B, Sereno MI (1999) Cortical surface-based analysis. I. Segmentation and surface reconstruction. *NeuroImage* 9:179–194
- Deichmann R, Gottfried JA, Hutton C, Turner R (2003) Optimized EPI for fMRI studies of the orbitofrontal cortex. *Neuroimage* 19:430–441
- Detre JA, Wang J (2002) Technical aspects and utility of fMRI using BOLD and ASL. *Clinical Neurophysiol* 113:621–634
- Donaldson DI, Petersen SE, Ollinger JM, Buckner RL (2001) Dissociating state and item components of recognition memory using fMRI. *NeuroImage* 13:129–142
- Draganski B, Gaser C, Busch V, Schuierer G, Bogdahn U, May A (2004) Neuroplasticity: changes in grey matter induced by training. *Nature* 427:311–312
- Ernst T, Hennig J (1994) Observation of a fast response in functional MR. *Magn Reson Med* 32:146–149
- Figley CR, Stroman PW (2011) The role(s) of astrocytes and astrocyte activity in neurometabolism, neurovascular coupling, and the production of functional neuroimaging signals. *Eur J Neurosci* 33:577–588
- Fischl B, Sereno MI, Dale AM (1999a) Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system. *Neuroimage* 9:195–207
- Fischl B, Sereno MI, Tootell RB, Dale AM (1999b) High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum Brain Mapp* 8:272–284
- Friston KJ (1998) Modes or models: a critique on independent component analysis for fMRI. *Trends Cogn Sci* 2:373–375
- Friston KJ (2005) Models of brain function in neuroimaging. *Annu Rev Psychol* 56:57–87
- Friston K (2011) Functional and effective connectivity: a review. *Brain Connectivity* 1:13–36

- Friston KJ, Holmes AP, Worsley KJ, Poline JP, Frith CD, Frackowiak RSJ (1994a) Statistical parametric maps in functional imaging: a general linear approach. *Hum Brain Mapp* 2:189–210
- Friston KJ, Tononi G, Reeke GN Jr, Sporns O, Edelman GM (1994b) Value-dependent selection in the brain: simulation in a synthetic neural model. *Neuroscience* 59:229–243
- Friston KJ, Price CJ, Fletcher P, Moore C, Frackowiak RS, Dolan RJ (1996) The trouble with cognitive subtraction. *Neuroimage* 4:97–104
- Friston KJ, Zarahn E, Josephs O, Henson RN, Dale AM (1999) Stochastic designs in event-related fMRI. *Neuroimage* 10:607–619
- Friston KJ, Glaser DE, Henson RN, Kiebel S, Phillips C, Ashburner J (2002) Classical and Bayesian inference in neuroimaging: applications. *Neuroimage* 16:484–512
- Gati JS, Menon RS, Ugurbil K, Rutt BK (1997) Experimental determination of the BOLD field strength dependence in vessels and tissue. *Magn Reson Med* 38:296–302, Official journal of the Society of Magnetic Resonance in Medicine/Society of Magnetic Resonance in Medicine
- Genovese CR, Lazar NA, Nichols T (2002) Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage* 15:870–878
- Gholipour A, Kehtarnavaz N, Briggs R, Devous M, Gopinath K (2007) Brain functional localization: a survey of image registration techniques. *IEEE Trans Med Imaging* 26:427–451
- Good CD, Johnsrude IS, Ashburner J, Henson RN, Friston KJ, Frackowiak RS (2001) A voxel-based morphometric study of ageing in 465 normal adult human brains. *Neuroimage* 14:21–36
- Gossel C, Auer DP, Fahrmeir L (2000) Dynamic models in fMRI. *Magn Reson Med* 43:72–81
- Harris JJ, Reynell C, Attwell D (2011) The physiology of developmental changes in BOLD functional imaging signals. *Dev Cogn Neurosci* 1:199–216
- Hartman D, Hlinka J, Palus M, Mantini D, Corbetta M (2011) The role of nonlinearity in computing graph-theoretical properties of resting-state functional magnetic resonance imaging brain networks. *Chaos* 21:013119
- Hoenig K, Kuhl CK, Scheef L (2005) Functional 3.0-T MR assessment of higher cognitive function: are there advantages over 1.5-T imaging? *Radiology* 234:860–868
- Hoge RD, Atkinson J, Gill B, Crelier GR, Marrett S, Pike GB (1999a) Investigation of BOLD signal dependence on cerebral blood flow and oxygen consumption: the deoxyhemoglobin dilution model. *Magn Reson Med* 42:849–863
- Hoge RD, Atkinson J, Gill B, Crelier GR, Marrett S, Pike GB (1999b) Linear coupling between cerebral blood flow and oxygen consumption in activated human cortex. *Proc Natl Acad Sci USA* 96:9403–9408
- Holmes AP, Blair RC, Watson JD, Ford I (1996) Nonparametric analysis of statistic images from functional mapping experiments. *J Cereb Blood Flow Metab* 16:7–22
- Hyvarinen A, Oja E (2000) Independent component analysis: algorithms and applications. *Neural Netw* 13:411–430
- Josephs O, Henson RN (1999) Event-related functional magnetic resonance imaging: modelling, inference and optimization. *Phil Trans R Soc Lond B Biol Sci* 354:1215–1228
- Kingsley PB (2006a) Introduction to diffusion tensor imaging mathematics: Part I. Tensors, rotations, and eigenvectors. *Concept Magn Reson A* 28A:101–122
- Kingsley PB (2006b) Introduction to diffusion tensor imaging mathematics: Part II. Anisotropy, diffusion-weighting factors, and gradient encoding schemes. *Concept Magn Reson A* 28A:123–154
- Kingsley PB (2006c) Introduction to diffusion tensor imaging mathematics: Part III. Tensor calculation, noise, simulations, and optimization. *Concept Magn Reson A* 28A:155–179
- Klein A, Andersson J, Ardekani BA, Ashburner J, Avants B, Chiang MC, Christensen GE, Collins DL, Gee J, Hellier P, Song JH, Jenkinson M, Lepage C, Rueckert D, Thompson P, Vercauteren T, Woods RP, Mann JJ, Parsey RV (2009) Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. *Neuroimage* 46:786–802
- Kruggel F, von Cramon DY (1999) Modeling the hemodynamic response in single-trial functional MRI experiments. *Magn Reson Med* 42:787–797

- Kwong KK (2011) Record of a single fMRI experiment in May of 1991. *Neuroimage*. doi:10.1016/j.neuroimage.2011.07.089 [Epub ahead of print]
- Li YO, Yang FG, Nguyen CT, Cooper SR, Lahue SC, Venugopal S, Mukherjee P (2011) Independent component analysis of DTI reveals multivariate microstructural correlations of white matter in the human brain. *Hum Brain Mapp*. doi:10.1002/hbm.21292 [Epub ahead of print]
- Liu G, Sobering G, Duyn J, Moonen CT (1993) A functional MRI technique combining principles of echo-shifting with a train of observations (PRESTO). *Magn Reson Med* 30:764–768
- Logothetis NK (2008) What we can do and what we cannot do with fMRI. *Nature* 453:869–878
- Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412:150–157
- Magri C, Logothetis NK, Panzeri S (2011) Investigating static nonlinearities in neurovascular coupling. *Magn Reson Imaging* 29:1358–1364
- Maguire EA, Gadian DG, Johnsrude IS, Good CD, Ashburner J, Frackowiak RS, Frith CD (2000) Navigation-related structural change in the hippocampi of taxi drivers. *Proc Natl Acad Sci USA* 97:4398–4403
- Mandeville JB, Marota JJ, Ayata C, Zaharchuk G, Moskowitz MA, Rosen BR, Weisskoff RM (1999) Evidence of a cerebrovascular postarteriole windkessel with delayed compliance. *J Cereb Blood Flow Metab* 19:679–689
- Mangia S, Giove F, Tkac I, Logothetis NK, Henry PG, Olman CA, Maraviglia B, Di Salle F, Ugurbil K (2009) Metabolic and hemodynamic events after changes in neuronal activity: current hypotheses, theoretical predictions and in vivo NMR experimental findings. *J Cereb Blood Flow Metab* 29:441–463
- McKeown MJ, Makeig S, Brown GG, Jung TP, Kindermann SS, Bell AJ, Sejnowski TJ (1998) Analysis of fMRI data by blind separation into independent spatial components. *Hum Brain Mapp* 6:160–188
- Menon RS, Kim SG (1999) Spatial and temporal limits in cognitive neuroimaging with fMRI. *Trends Cogn Sci* 3:207–216
- Menon RS, Ogawa S, Hu X, Strupp JP, Anderson P, Ugurbil K (1995) BOLD based functional MRI at 4 Tesla includes a capillary bed contribution: echo-planar imaging correlates with previous optical imaging using intrinsic signals. *Magn Reson Med* 33:453–459
- Menon RS, Luknowsky DC, Gati JS (1998) Mental chronometry using latency-resolved functional MRI. *Proc Natl Acad Sci USA* 95:10902–10907
- Miezin FM, Maccotta L, Ollinger JM, Petersen SE, Buckner RL (2000) Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *Neuroimage* 11:735–759
- Muthukumaraswamy SD, Evans CJ, Edden RA, Wise RG, Singh KD (2011) Individual variability in the shape and amplitude of the BOLD-HRF correlates with endogenous GABAergic inhibition. *Hum Brain Mapp* 33:455–465
- Neggers SF, Hermans EJ, Ramsey NF (2008) Enhanced sensitivity with fast three-dimensional blood-oxygen-level-dependent functional MRI: comparison of SENSE-PRESTO and 2D-EPI at 3 T. *NMR Biomed* 21:663–676
- Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Hum Brain Mapp* 15:1–25
- Nirxko AC (2003) Nitric oxide-an endogenous contrast agent contributing to “the elusive initial dip?”. *NeuroImage* 20:611–612
- Ogawa S, Lee TM, Kay AR, Tank DW (1990) Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc Natl Acad Sci USA* 87:9868–9872
- Ogawa S, Menon RS, Tank DW, Kim SG, Merkle H, Ellermann JM, Ugurbil K (1993) Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model. *Biophys J* 64:803–812
- Osterbauer RA, Wilson JL, Calvert GA, Jezzard P (2006) Physical and physiological consequences of passive intra-oral shimming. *Neuroimage* 29:245–253

- Parker GJ, Alexander DC (2003) Probabilistic Monte Carlo based mapping of cerebral connections utilising whole-brain crossing fibre information. *Inf Process Med Imaging* 18:684–695
- Parker GJ, Haroon HA, Wheeler-Kingshott CA (2003) A framework for a streamline-based probabilistic index of connectivity (PICO) using a structural interpretation of MRI diffusion measurements. *J Magn Reson Imaging: JMRI* 18:242–254
- Pauling L, Coryell CD (1936) The magnetic properties and structure of hemoglobin, oxyhemoglobin and carbonmonoxyhemoglobin. *Proc Natl Acad Sci USA* 22:210–216
- Petersson KM, Nichols TE, Poline JB, Holmes AP (1999a) Statistical limitations in functional neuroimaging. I. Non-inferential methods and statistical models. *Phil Trans R Soc Lond B Biol Sci* 354:1239–1260
- Petersson KM, Nichols TE, Poline JB, Holmes AP (1999b) Statistical limitations in functional neuroimaging. II. Signal detection and statistical inference. *Philos Trans R Soc Lond B Biol Sci* 354:1261–1281
- Poline JB, Worsley KJ, Evans AC, Friston KJ (1997) Combining spatial extent and peak intensity to test for activations in functional imaging. *Neuroimage* 5:83–96
- Roche A (2011) A four-dimensional registration algorithm with application to joint correction of motion and slice timing in fMRI. *IEEE Trans Med Imaging* 30:1546–1554
- Rubinov M, Sporns O (2010) Complex network measures of brain connectivity: uses and interpretations. *Neuroimage* 52:1059–1069
- Savoy RL (2005) Experimental design in brain activation MRI: cautionary tales. *Brain Res Bull* 67:361–367
- Savoy RL, O'Craven KM, Weisskoff RM, Davis TL, Baker J, Rosen BR (1994) Exploring the temporal boundaries of fMRI: measuring responses to very brief stimuli. 24th. Annual Meeting of the Society of Neuroscience, Miami, p 1264
- Schmiedeskamp H, Newbould RD, Pisani LJ, Skare S, Glover GH, Pruessmann KP, Bammer R (2010) Improvements in parallel imaging accelerated functional MRI using multiecho echo-planar imaging. *Magn Reson Med* 63:959–969
- Segonne F, Dale AM, Busa E, Glessner M, Salat D, Hahn HK, Fischl B (2004) A hybrid approach to the skull stripping problem in MRI. *Neuroimage* 22:1060–1075
- Skudlarski P, Jagannathan K, Calhoun VD, Hampson M, Skudlarska BA, Pearlson G (2008) Measuring brain connectivity: diffusion tensor imaging validates resting state temporal correlations. *Neuroimage* 43:554–561
- Sladky R, Friston KJ, Trostl J, Cunnington R, Moser E, Windischberger C (2011) Slice-timing effects and their correction in functional MRI. *NeuroImage* 58:588–594
- Smith SM (2002) Fast robust automated brain extraction. *Hum Brain Mapp* 17:143–155
- Smith SM, Jenkinson M, Johansen-Berg H, Rueckert D, Nichols TE, Mackay CE, Watkins KE, Ciccarelli O, Cader MZ, Matthews PM, Behrens TE (2006) Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *Neuroimage* 31:1487–1505
- Smith SM, Miller KL, Salimi-Khorshidi G, Webster M, Beckmann CF, Nichols TE, Ramsey JD, Woolrich MW (2011) Network modelling methods for FMRI. *Neuroimage* 54:875–891
- Sporns O (2011) *Networks of the brain*. MIT Press, Cambridge, MA
- Stephan KE, Friston KJ (2010) Analyzing effective connectivity with fMRI. *Wiley interdisciplinary reviews. Cogn Sci* 1:446–459
- Stone JV (2002) Independent component analysis: an introduction. *Trends Cogn Sci* 6:59–64
- Takao H, Hayashi N, Ohtomo K (2011) Effect of scanner in longitudinal studies of brain volume changes. *J Magn Reson Imaging: JMRI* 34:438–444
- Talairach J, Tournoux P (1988) *Co-planar stereotaxic Atlas of the human brain*. Thieme, Stuttgart
- Teipel SJ, Bokde AL, Meindl T, Amaro E Jr, Soldner J, Reiser MF, Herpertz SC, Moller HJ, Hampel H (2010) White matter microstructure underlying default mode network connectivity in the human brain. *Neuroimage* 49:2021–2032
- Tepest R, Wang L, Csernansky JG, Neubert P, Heun R, Scheef L, Jessen F (2008) Hippocampal surface analysis in subjective memory impairment, mild cognitive impairment and Alzheimer's dementia. *Dement Geriatr Cogn Disord* 26:323–329

- Thulborn KR, Waterton JC, Matthews PM, Radda GK (1982) Oxygenation dependence of the transverse relaxation time of water protons in whole blood at high field. *Biochim Biophys Acta* 714:265–270
- Tian P, Teng IC, May LD, Kurz R, Lu K, Scadeng M, Hillman EM, De Crespigny AJ, D'Arceuil HE, Mandeville JB, Marota JJ, Rosen BR, Liu TT, Boas DA, Buxton RB, Dale AM, Devor A (2010) Cortical depth-specific microvascular dilation underlies laminar differences in blood oxygenation level-dependent functional MRI signal. *Proc Natl Acad Sci USA* 107:15246–15251
- Triantafyllou C, Hoge RD, Krueger G, Wiggins CJ, Potthast A, Wiggins GC, Wald LL (2005) Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *Neuroimage* 26:243–250
- Valdes-Sosa PA, Sanchez-Bornot JM, Sotero RC, Iturria-Medina Y, Aleman-Gomez Y, Bosch-Bayard J, Carbonell F, Ozaki T (2009) Model driven EEG/fMRI fusion of brain oscillations. *Hum Brain Mapp* 30:2701–2721
- van den Heuvel MP, Hulshoff Pol HE (2010) Exploring the brain network: a review on resting-state fMRI functional connectivity. *Eur Neuropsychopharmacol* 20:519–534
- van der Zwaag W, Francis S, Head K, Peters A, Gowland P, Morris P, Bowtell R (2009) fMRI at 1.5, 3 and 7 T: characterising BOLD signal changes. *Neuroimage* 47:1425–1434
- Vanzetta I, Grinvald A (2001) Evidence and lack of evidence for the initial dip in the anesthetized rat: implications for human functional brain imaging. *Neuroimage* 13:959–967
- Wager TD, Nichols TE (2003) Optimization of experimental design in fMRI: a general framework using a genetic algorithm. *Neuroimage* 18:293–309
- Wagner AD, Schacter DL, Rotte M, Koutstaal W, Maril A, Dale AM, Rosen BR, Buckner RL (1998) Building memories: remembering and forgetting of verbal experiences as predicted by brain activity. *Science* 281:1188–1191
- Wang J, Aguirre GK, Kimberg DY, Roc AC, Li L, Detre JA (2003) Arterial spin labeling perfusion fMRI with very low task frequency. *Magn Reson Med* 49:796–802
- Watts DJ, Strogatz SH (1998) Collective dynamics of ‘small-world’ networks. *Nature* 393:440–442
- Weiskopf N, Hutton C, Josephs O, Turner R, Deichmann R (2007) Optimized EPI for fMRI studies of the orbitofrontal cortex: compensation of susceptibility-induced gradients in the readout direction. *MAGMA* 20:39–49
- Wilson JL, Jezzard P (2003) Utilization of an intra-oral diamagnetic passive shim in functional MRI of the inferior frontal cortex. *Magn Reson Med* 50:1089–1094
- Woods RP, Mazziotta JC, Cherry SR (1993) MRI-PET registration with automated algorithm. *J Comput Assist Tomogr* 17:536–546
- Woolrich MW, Behrens TE, Beckmann CF, Jenkinson M, Smith SM (2004a) Multilevel linear modelling for FMRI group analysis using Bayesian inference. *Neuroimage* 21:1732–1747
- Woolrich MW, Behrens TE, Smith SM (2004b) Constrained linear basis sets for HRF modelling using Variational Bayes. *Neuroimage* 21:1748–1761
- 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
- Wright IC, McGuire PK, Poline JB, Travers JM, Murray RM, Frith CD, Frackowiak RS, Friston KJ (1995) A voxel-based method for the statistical analysis of gray and white matter density applied to schizophrenia. *Neuroimage* 2:244–252
- Yacoub E, Shmuel A, Pfeuffer J, Van De Moortele PF, Adriany G, Andersen P, Vaughan JT, Merkle H, Ugurbil K, Hu X (2001a) Imaging brain function in humans at 7 Tesla. *Magn Reson Med* 45:588–594
- Yacoub E, Shmuel A, Pfeuffer J, Van De Moortele PF, Adriany G, Ugurbil K, Hu X (2001b) Investigation of the initial dip in fMRI at 7 Tesla. *NMR Biomed* 14:408–412
- Zarahn E, Aguirre G, D'Esposito M (1997) A trial-based experimental design for fMRI. *Neuroimage* 6:122–138
- Zeineh MM, Engel SA, Thompson PM, Bookheimer SY (2001) Unfolding the human hippocampus with high resolution structural and functional MRI. *Anat Rec* 265:111–120