

Revealing Brain Activity and White Matter Structure Using Functional and Diffusion-Weighted Magnetic Resonance Imaging

Rainer Goebel

Contents

1	Physical Principles of MRI	14
1.1	Spin Excitation and Signal Reception.....	15
1.2	Image Reconstruction	18
2	Physiological Principles of fMRI	21
2.1	Neurovascular Coupling.....	21
2.2	The BOLD Effect.....	22
2.3	The BOLD Hemodynamic Response.....	22
2.4	Limits of Spatial and Temporal Resolution	24
3	FMRI Data Analysis	24
3.1	Block- and Event-Related Designs	25
3.2	Basic Analysis Steps	27
3.3	Statistical Analysis of Functional Data.....	31
3.4	Integration of Anatomical and Functional Data.....	44
3.5	Group Analysis of Functional Data Sets.....	45
3.6	Selected Advanced Data Analysis Methods	49
4	Functional Connectivity and Resting-State Networks	51
4.1	Functional and Effective Connectivity.....	51
4.2	Resting-State Networks.....	51
5	Diffusion-Weighted MRI and Tractography	53
5.1	Diffusion Tensor Imaging	53
5.2	Validation and Improvements	55
5.3	Applications	55
	References	57

Abstracts

Magnetic resonance imaging (MRI) is based on the magnetic *excitation* of body tissue and the reception of returned electromagnetic signals from the body. Excitation induces phase-locked *precession* of protons with a frequency proportional to the strength of the surrounding magnetic field as described by the *Larmor equation*. This fact can be exploited for *spatial encoding* by applying magnetic field *gradients* along spatial dimensions on top of the strong static magnetic field of the scanner. The obtained frequency-encoded information for each slice is accumulated in two-dimensional *k space*. The k space data can be transformed into *image space* by *Fourier analysis*.

Functional MRI (fMRI) allows *localizing brain function* since increased local neuronal activity leads to a surprisingly strong increase in local blood flow, which itself results in measurable increases in local magnetic field homogeneity. Increased local blood flow delivers chemical energy (glucose and oxygen) to the neurons. The temporary increase and decrease of local blood flow, triggered by increased neuronal activity, is called the *hemodynamic response* starting 2–4 s after stimulus onset. Increased local blood flow results in an *oversupply* of *oxygenated hemoglobin* in the vicinity of increased neuronal activity. The oversupply flushes *deoxygenated hemoglobin* from the capillaries and the downstream venules. Deoxygenated hemoglobin is paramagnetic reducing the homogeneity of the local magnetic field resulting in a weaker MRI signal than would be measurable without it. Oxygenated hemoglobin is diamagnetic and does not strongly reduce field homogeneity. Since the increased local blood flow replaces deoxygenated hemoglobin with oxygenated hemoglobin, local field homogeneity increases, leading to a stronger MRI signal as compared to a nonactivated state. Measured functional brain images thus reflect neuronal activity changes as blood oxygenation level-dependent (*BOLD*) contrast.

Functional images are acquired using the fast *echo-planar imaging* (EPI) pulse sequence allowing acquisition of a 64×64 image matrix in less than 100 ms. To sample

R. Goebel
Department of Cognitive Neuroscience, Maastricht University,
Oxfordlaan 55, 6229 EV Maastricht, The Netherlands
e-mail: r.goebel@maastrichtuniversity.nl

signal changes over time, a set of slices typically covering the whole brain is *measured repeatedly*. Activation of neurons results in a BOLD signal increase of only about 1–5 % and lie buried within strong physical and physiological noise fluctuations of similar size. Proper *preprocessing* steps, including 3D motion correction and removal of drifts, reduce the effect of artifacts increasing the *signal-to-noise ratio* (SNR). In order to reliably *detect* stimulus-related effects, proper *statistical data analysis* is performed. In order to *estimate* response profiles, condition-related time course episodes may be averaged in various *regions-of-interest* (ROIs). The core statistical tool in fMRI data analysis is the *general linear model* (GLM) that allows the analysis of *blocked* and *event-related experimental designs*. To run a GLM, a *design matrix* (model) has to be constructed containing reference functions (*predictors*, model time courses) for all effects of interest (conditions) as well as confounds. The GLM fits the created model to the data independently for each voxel's data (time course) providing a set of beta values estimating the effects of each condition. These beta values are compared with each other using *contrasts* resulting in a statistical value at each voxel. The statistical values of all voxels form a three-dimensional *statistical map*. To protect against wrongly declaring voxels as significant, statistical maps are *thresholded* properly by taking into account the *multiple comparison problem*. This problem is caused by the large number of independently performed statistical tests (one for each voxel).

In recent years, *parallel imaging techniques* have been developed, which allow acquiring MRI data simultaneously with two or more receiver coils. Parallel imaging can be used to increase temporal or spatial resolution. It also helps to reduce EPI imaging artifacts, such as geometrical distortions and signal dropouts in regions of different neighboring tissue types.

MRI has not only revolutionized functional brain imaging targeting gray matter neuronal activity but also enabled insights into the human white matter structure using diffusion-weighted magnetic resonance imaging. With proper measurement and modeling schemes including diffusion tensor imaging (DTI), major long-range fiber tracts can be reconstructed using computational tractography providing important information to guide neurosurgical procedures potentially reducing the risk of lesioning important fiber tracts.

Since its invention in the early 1990s, functional magnetic resonance imaging (fMRI) has rapidly assumed a leading role among the techniques used to localize brain activity. The spatial and temporal resolution provided by state-of-the-art MR technology and its non-invasive character, which allows multiple studies of the same subject, are some of the main advantages of fMRI over the other functional neuroimaging

techniques that are based on changes in blood flow and cortical metabolism (e.g., positron-emission tomography, PET). fMRI is based on the discovery of Ogawa et al. (1990), that *magnetic resonance imaging* (MRI, also called nuclear magnetic resonance imaging) can be used in a way that allows obtaining signals depending on the level of blood oxygenation. The measured signal is therefore also called “BOLD” signal (BOLD=*blood oxygenation level-dependent*). Since locally increased neuronal activity leads to increased local blood flow, which again changes local blood oxygenation, fMRI allows indirect measurements of neuronal activity changes. With appropriate data analysis and visualization methods, these BOLD measurements allow drawing conclusions about the localization and dynamics of brain function.

This chapter describes the basic principles and methodology of functional and diffusion-weighted MRI. After a description of the physical principles of MRI at a conceptual level, the physiology of the blood oxygenation level-dependent (BOLD) contrast mechanism is described. The subsequent, major part of the chapter provides an introduction to the current strategies of statistical image analysis techniques with a focus on the analysis of single-subject data because of its relevance for pre-surgical mapping of human brain function. This is followed by a description of functional connectivity focusing on the analysis of resting state fMRI data. Finally, principles of diffusion-weighted MRI measurements are described including diffusion tensor imaging, which is the most common acquisition and modeling approach in clinical MRI.

1 Physical Principles of MRI

Magnetic resonance imaging allows visualizing both anatomical and functional data of the human brain. This section shortly describes the main concepts of the physical principles of MRI. More detailed descriptions of the physical basis of MRI are available in several introductory texts, for example, Huettel et al. (2004), Bandettini et al. (2000), Brown and Semelka (1999), NessAiver (1997), and Schild (1990).

A typical whole-body MR scanner has a hollow bore (tube) about 1 m across. Inside of that bore a cylinder is placed containing the primary magnet producing a very strong *static, homogeneous magnetic field* (B_0). Today, nearly all scanners create the magnetic field with superconducting electromagnets whose wires are cooled by cryogenics (e.g., liquid helium). Most standard clinical scanners used to image the human brain possess a magnetic field strength of 1.5 T, which is 30,000 times the strength of the Earth's magnetic field (1 T = 10,000 G). In recent years, installation of tomographs with 3.0 T has become common in major hospitals and research centers. In a few research labs, the human brain is imaged at ultrahigh fields such as 7 and 9.4 T. At higher field strengths it gets increasingly difficult to create a homogeneous magnetic field, which is necessary for accurate spatial

decoding of the raw measurement data. Since homogenous fields are easier to create for scanners with small bores, scanners with higher magnetic fields (10–20 T) are currently only available for animal use. Besides the main magnet, additional coils are located inside the cylinder including shimming coils, gradient coils, and a radio frequency (RF) coil. The shimming coils are used to shape the magnetic field increasing its homogeneity. The gradient coils are used to temporarily change the magnetic field linearly along any direction which is essential for spatial localization (see below). The RF coil is used to send radio frequency pulses into the subject.

In a typical brain scanning session, a subject or patient in supine position is slowly moved into the scanner bore using a maneuverable table. Scanning of anatomical and functional images is managed from a terminal in a control room by specifying slice positions and by running appropriate MRI *pulse sequences*. The control room usually has a window behind the computer terminal, which allows looking into the scanner room. Before the subject is moved into the scanner, the head is placed in a small replaceable coil, called the *head coil*. This coil surrounds the head and is used to send radio frequency pulses into the subject as well as to receive electromagnetic echoes. When receive-only head coils are used, the radio frequency pulses are provided by the RF coil in the cylinder of the scanner. The head coil is an example of a *volume coil*, which is designed such that the sensitive volume (e.g., brain) experiences a fairly uniform RF field. Surface coils are receive-only RF coils that are placed directly upon the surface of the anatomy to be imaged. They provide very high signal-to-noise in their immediate vicinity, but recorded images suffer from extreme nonuniformity because the obtained signal intensity drops rapidly with distance and approaches zero about one coil diameter away from the coil. Phased array coils

are an attempt to combine the positive properties of volume and surface coils by combining images from two or more surface coils to produce a single image (see Sect. 1.2.6).

The physical principles of MRI are the same for anatomical and functional imaging. What makes functional imaging special is described in Sect. 2. The operation of MRI can be described in two major themes. The first theme refers to the excitation and recording of electromagnetic signals reflecting the properties of the measured object. The second theme refers to the construction of two- and three-dimensional images reflecting how the measured object properties vary across space.

1.1 Spin Excitation and Signal Reception

Magnetic resonance imaging is based on the magnetic excitation of body tissue and the recording of returned electromagnetic signals from the body. All nuclei with an odd number of protons are magnetically excitable. The atom of choice for MRI is ^1H , the most common isotope of hydrogen having a nucleus with only one proton. Hydrogen protons are ideally suited for MRI because they are abundant in human tissue and possess particularly favorable magnetic properties. Water is the largest source of protons in the body followed by fat. Protons have magnetic properties because they possess a *spin*: they rotate like a spin top around their own axes inducing a small directed magnetic field. In a normal environment, the magnetic fields of the spins in the human body are oriented randomly and, thus, cancel each other out. If, however, the body of a subject is placed in the strong static magnetic field of a MRI tomography (called \mathbf{B}_0), the spins orient themselves in line with that field, either parallel or antiparallel (Fig. 1). Since a slightly larger proportion of spins aligns parallel to the

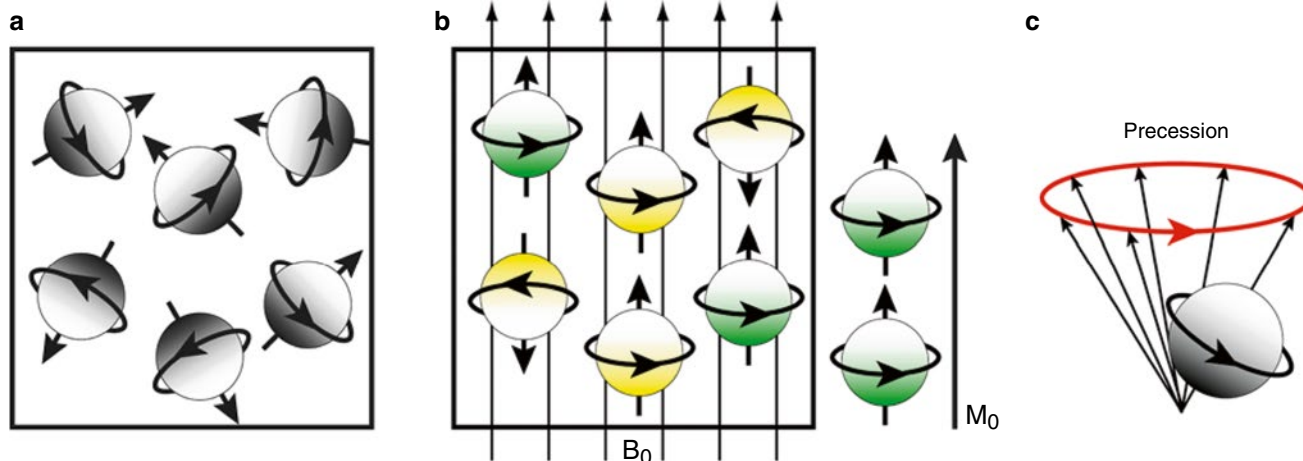


Fig. 1 Spinning protons are little magnets because of the spin property. (a) Without an external magnetic field, the directions of the spins are randomly distributed. (b) When placed within a large magnetic field, the spins align either with the field (parallel) or against the field (antiparallel). A slight excess of spins aligns with the external magnetic

field resulting in a net magnetic field parallel to the external magnetic field. (c) A spin is actually not aligning its axis of rotation with the external magnetic field as shown in (a) and (b) but rotates around the direction of the field. This motion is called precession




Fig. 2 Spins in the lower energy state can be excited by an electromagnetic pulse at the resonance frequency ω_0 forcing the spins absorbing the transmitted energy to precess in phase. (a) As an effect of excitation, the net magnetic field \mathbf{M}_0 (blue vector) smoothly tips down toward the x - y plane. The longitudinal component \mathbf{M}_z (green vector) decreases over time while the transverse component \mathbf{M}_{xy} (red vector)

increases. This view assumes that the observer is moving with the precessing protons (rotating frame of reference). (b) Viewed from outside (laboratory frame of reference), the net magnetization vector rotates with angular velocity ω_0 given by the Larmor equation. The rotating magnetic field in the x - y plane emits radiofrequency waves, which can be measured by a receiver coil

scanner magnetic field, the body gets magnetized. The excess number of spins aligned with the external magnetic field is proportional to the strength of the external magnetic field and is in the order of 10^{15} spins at 1.5 T in a $2 \times 2 \times 2$ mm volume of water. The total magnetic field of the excess spins is called \mathbf{M}_0 . Just as a spinning top wobbles about its axis, the spinning protons wobble, or *precess*, about the axis of the external \mathbf{B}_0 field (Fig. 1c). The precession frequency of the protons depends on the strength of the surrounding magnetic field. More precisely, the precession frequency ω is directly proportional to the strength of the external magnetic field and is defined by the *Larmor equation*:

$$\omega_0 = \gamma \mathbf{B}_0$$

The symbol ω_0 is known as the precessional, Larmor, or *resonance* frequency. The symbol γ refers to the *gyromagnetic ratio*, which is a constant unique to every atom. For hydrogen protons, $\gamma = 42.56$ MHz per Tesla. At the magnetic field strength of a 3 T scanner, the precession frequency of hydrogen protons is thus 128 MHz.

If an applied electromagnetic pulse has the same frequency as the proton's precession frequency, then the protons get "excited" by absorbing the transmitted energy. This important principle is called *resonance* and gives the method "magnetic resonance imaging" its name. Since the precession frequency is in the range of radio frequency waves, the applied electromagnetic pulse is also called a radio frequency (RF) pulse. As an effect of excitation, spins flip from the parallel (lower-energy) state to the antiparallel (higher-energy) state. The RF pulse furthermore lets the excited protons precess *in phase*. As a result, the magnetization vector \mathbf{M}_0 moves down toward the x - y plane (Fig. 2). The x - y plane is perpendicular to the static magnetic field and is also referred to as the *transverse plane*. The angle, α , of rotation

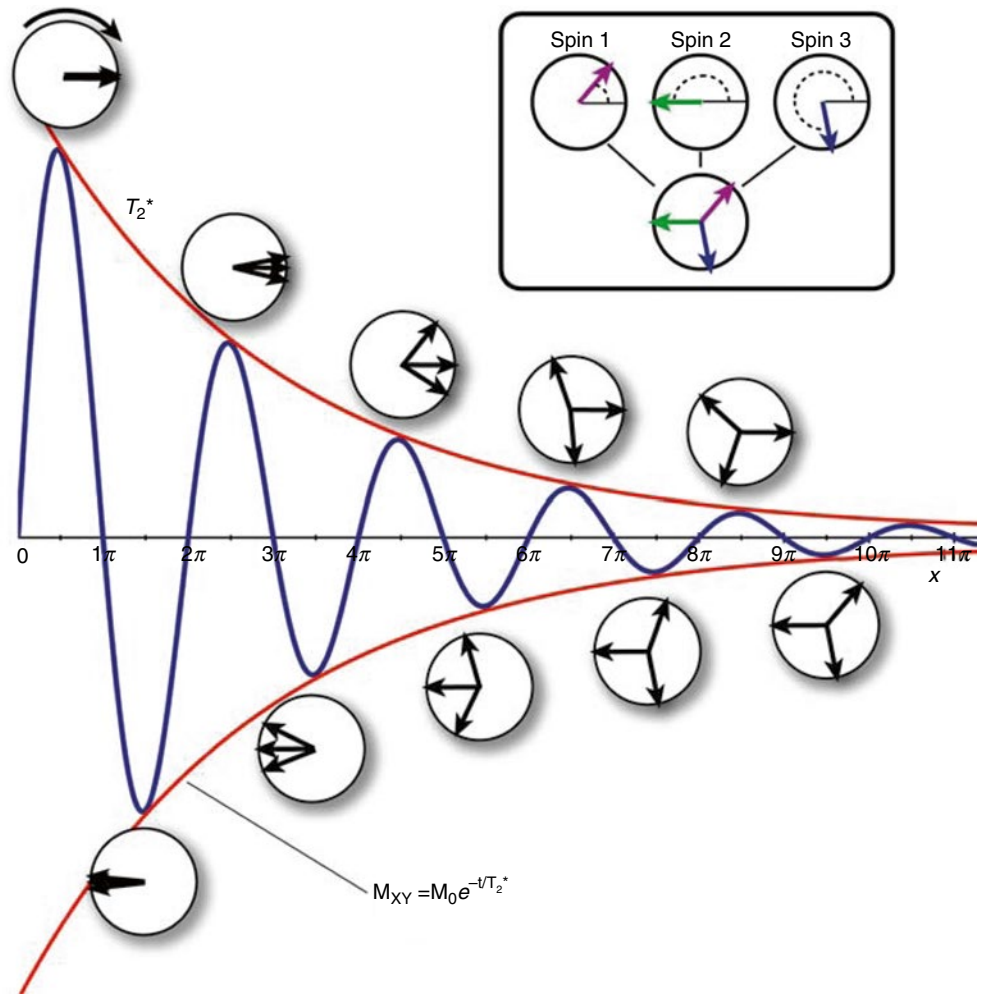
toward the x - y plane is a function of the strength and duration of the RF pulse. If $\alpha = 90^\circ$, the magnetization vector is completely moved into the x - y plane with an equal amount of spins aligned parallel and antiparallel (Fig. 2b). Since the protons precess in phase, that is, they point in the same direction within the x - y plane, the magnetic fields of the spins add up to form a net magnetic field \mathbf{M}_{xy} in the x - y plane. This *transversal component* of the rotating electromagnetic field can be measured (received) in the receiver coil (antenna) because it induces a detectable current flow.

The established inphase precession is, however, not stable after the RF transmitter is turned off. Because of interactions between the magnetic fields of the protons, the transverse magnetization decays within a few tenths of milliseconds. These spin-spin interactions lead to slightly different local magnetic field strengths and, thus, to slightly different precession frequencies leading to phase shifts between the precessing spins (*dephasing*). The dephasing process is also called *transversal relaxation*. It progresses initially rapidly but slows down over time following an exponential function with time constant T_2 with values in the range of 30–150 ms. Due to magnetic field inhomogeneities in the static magnetic field and in physiological tissue, the spins get out of phase actually faster than T_2 , and therefore the measured raw signal in the receiver coil, the *free induction decay* (FID), decays with a shorter time constant T_2^* (Fig. 3):

$$\mathbf{M}_{xy} = \mathbf{M}_0 e^{-t/T_2^*}$$

The fact that local field inhomogeneities lead to different precession frequencies increasing the speed of dephasing is an important observation for functional MRI because local field inhomogeneities also depend on the local physiological state, especially the state of local blood oxygenation, which

Fig. 3 The signal amplitude (red curve) of the measured raw MR signal, the free induction decay (FID), decays exponentially with time constant T_2^* . The raw signal itself is oscillating at the resonance frequency (blue curve). The signal is lost due to dephasing as indicated by the phase coherency plots (circles) with three representative, superimposed spins (see inset). The amplitude of the signal at any moment in time is determined by the sum of the spin vectors. When the spins are all in phase (left side), the maximum signal is obtained, that is, the vector sum equals M_0 . When the spins are completely out of phase (right side), the signal is completely lost, that is, the sum of the spin vectors equals zero



itself depends on the state of local neuronal activity. Measurements of changing local magnetic field inhomogeneities (T_2^* parameter), thus, provide indirect measurements of local neuronal activity.

The speed of spin dephasing is determined by random effects as well as by fixed effects due to magnetic field inhomogeneities. The dephasing effect of constant magnetic field inhomogeneities can be reversed by the application of a 180° RF pulse. A time duration of $t = \tau$ is allowed to elapse while the spins go out of phase. Then a 180° RF pulse is applied flipping the dephased spin vectors about the X' or Y' axis in the rotating frame of reference. As an effect of the pulse, the order of the spins is reversed (Fig. 4). At the *echo time* $TE = 2\tau$, the vectors are back in phase producing a large signal, the *spin echo*. This process is similar to a race situation in which participants run with different (but constant) speed. At time τ they get a signal (“ 180° pulse”) to turn around and go back; assuming they continue in the same speed, they will all arrive at the starting line at the same time (2τ).

The amplitude of the obtained spin echo will be smaller than the amplitude during the FID because part of the signal

is inevitably lost due to random spin-spin interactions (T_2 decay). As soon as the spins are all back in phase at the echo time, they immediately start to go out of phase again. An additional 180° RF pulse will generate a second echo (Fig. 4). This process can be continued as long as enough signal is available. By setting the time of the 180° pulse, the amplitude of the T_2 signal can, thus, be assessed at any moment in time.

Besides dephasing, the spins reorient themselves with the direction of the strong static magnetic field of the scanner since the excited spins slowly go back into low-energy state realigning with the external magnetic field. This reorientation process is called *longitudinal relaxation* and progresses slower as the dephasing process. The increase (recovery) of the longitudinal component M_z follows an exponential function with time constant T_1 with values in the range of 300–2,000 ms:

$$M_z = M_0 (1 - e^{-t/T_1})$$

Note that the absorbed RF energy is not only released in a way that it can be detected outside the body as RF waves but part of the energy is given to the surrounding tissue, called the *lattice*.




Fig. 4 The effect of constant magnetic field inhomogeneities can be reversed by application of a 180° RF pulse, which flips the dephased vectors about the X' axis. This is indicated in the upper row with three spin vectors, one precessing at the resonance frequency (*green vector*), one precessing slightly faster (*violet vector*), and one precessing slightly slower (*blue vector*) leading to dephasing. The 180° RF pulse reverses the order of the spins but not the direction of rotation. The

faster spin now runs behind catching up over time, while the slower spin runs ahead slowly falling back. At time TE (echo time), the vectors are back in phase producing a large signal, the spin echo. A second 180° RF pulse will generate a second echo (*right side*). The maximum amplitude of the echoes gets smaller over time because signal is inevitably lost due to random spin-spin interactions (T_2 decay, *red curve*)

The spin-lattice interactions determine the speed of T_1 recovery, which is unique to every tissue. Tissue-specific T_1 and T_2 values enable MRI to differentiate between different types of tissue when using properly designed MRI pulse sequences.

1.2 Image Reconstruction

The described principles of magnetic resonance do not explain how one can obtain images of the brain. This requires attributing components of the signal to those positions in space from which they originated. Although not identical for all measurement sequences, the principles for localizing signal sources typically contain the combined application of three fundamental techniques: *selective excitation of a slice*, *frequency encoding*, and *phase encoding*. Each of these steps allows localizing the source of the signal with respect to one spatial dimension. Paul C. Lauterbur and Peter Mansfield were awarded the 2003 Nobel Prize in Medicine for their discovery that *magnetic field gradients* can be used for spatial encoding. The gradient coils of the MRI scanner allow adding a magnetic field to the static magnetic field, which causes the field strength to vary linearly with distance from the center of the magnet. According to the Larmor equation, spins on one side are exposed to a higher magnetic field and

precess faster while spins on the other side are exposed to a lower magnetic field and precess slower than spins in the center (Fig. 5b).

1.2.1 Selective Slice Excitation

A magnetic field gradient is used to select a slice of the imaged object (*slice selection gradient*). Since spins precess with different frequencies along a gradient, protons can be excited selectively: An applied electromagnetic pulse of a certain frequency band will excite only those protons along the gradient precessing at the same frequency band. Spins outside that range will precess at different frequencies and will, thus, not absorb the transmitted RF energy. The selectively excited protons are located in a slice oriented perpendicular to the gradient direction. A gradient along the z -axis will result in an axial slice, a gradient along the x -axis in a sagittal slice, and a gradient along the y -axis in a coronal slice. Oblique slices can be obtained by applying two or three gradients simultaneously. The position and thickness of the selected slice depend on the slope of the applied gradient and the frequency band of the applied RF pulse. After selective slice excitation, the measured echo will be restricted to a compound signal from the excited protons within the slice. For subsequent spatial encoding, the slice selection gradient is turned off.




Fig. 5 Assume that eight glasses with different amounts of water are placed in the MRI scanner along the x -axis and that a single, thick slice containing all glasses has been excited. **(a)** In the absence of any gradients, all of the excited protons from all glasses are spinning at the same frequency. The received signal also oscillates at that frequency and its amplitude reflects the sum of excited water protons of all glasses. Since all protons precess at the same frequency, the Fourier transform cannot be used to identify signals from different spatial positions along the x -axis. **(b)** If a gradient is applied in the x direction, the spins will

precess at frequencies that depend upon their position along the gradient. Spatial information is now frequency encoded: The strength of the signal at each frequency is directly related to the number of excited protons from the respective glass of water. The obtained composite time-domain signal is the sum of these frequencies. The Fourier transform can now be used to determine the strength of the signal at each frequency. Since frequencies encode different spatial positions, an “image” of eight pixels can be formed. The gray values of these pixels reflect the relative amount of water in the different glasses

1.2.2 Frequency Encoding

While receiving the signal (FID or echo) from the excited slice, a magnetic field gradient can be applied along one of the two remaining spatial dimensions. This second gradient, running along one dimension of the excited slice, is called *frequency-encoding gradient*. Note that this gradient is not used to selectively excite protons but to encode a spatial dimension for those protons already excited in the slice. Due to the applied gradient, the protons within the slice precess with different frequencies along the respective dimension allowing differentiating spatial positions in the received signal (Fig. 5). The frequency-encoding gradient is also called *readout gradient* since it is tuned on during reception of the signal from the protons. The strength of the signal at each frequency is directly related to the strength of the signal at the encoded spatial position. The measured composite time-domain signal consists of the sum of all frequency responses. The Fourier transform (FT) can be used to get from the composite signal the strength of the signal at each frequency (amplitude and phase information). Since space has been frequency encoded, the FT provides the strength of the signal at different spatial positions. The obtained frequency-specific information can thus be used to form a spatial image (Fig. 5b). In such an image, the gray level is used to represent the strength of the signal at each picture element (pixel).

1.2.3 Phase Encoding

A further encoding step is required in order to be able to also separate signal components originating from different positions along the second dimension in the imaging plane. This is achieved by briefly adding another gradient to the static magnetic field oriented along the remaining (third) spatial dimension before receiving an echo. This third magnetic field gradient is called *phase-encoding gradient*. While the frequency-encoding gradient is turned on during reception of the signal, the phase-encoding (PE) gradient is *turned off* just before receiving the echo and is, thus, not (permanently) changing the frequency at different spatial positions. This is necessary since frequency-encoding gradients in two dimensions would result in ambiguous spatial encoding in a similar way as the same number (e.g., 6) can be obtained in many different ways by the sum of two numbers (e.g., $2+4$, $3+3$, $5+1$). Prior to readout, the brief duration of the phase-encoding gradient results in a short moment of different precession frequencies within each row of the slice. After turning off the phase-encoding gradient, the protons within each row precess again with the same frequency but they will now precess with a systematic *phase shift* along the positions within each row. The amount of phase shift depends on the position of a proton along the encoded second image dimension. Through proper combination of frequency encoding in one dimension and phase encoding in the other dimension,

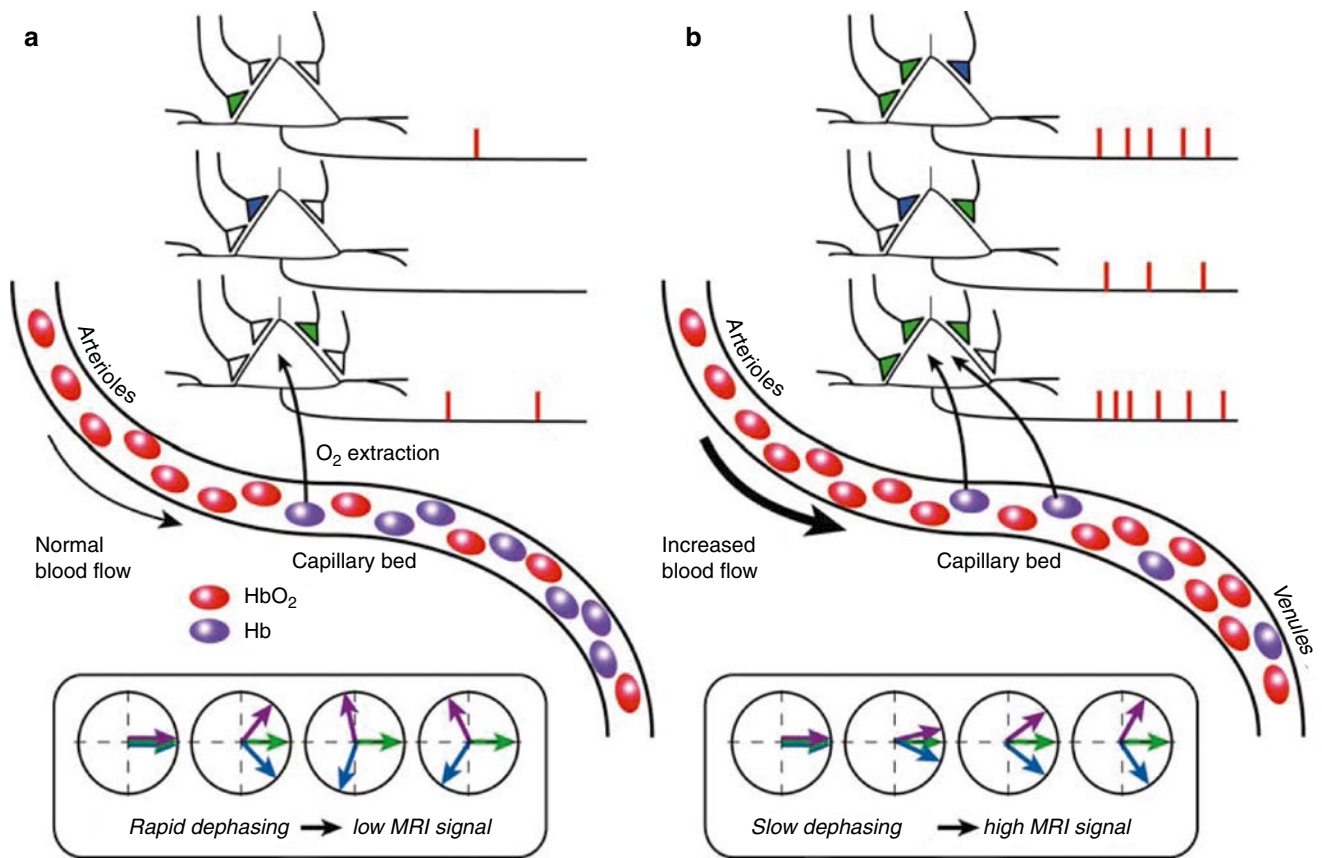


Fig. 6 From neural activity to BOLD MRI responses. **(a)** If a cortical region is in baseline mode, neural activity – including synaptic signal integration and spike generation – is low (*upper part*). Cerebral blood flow (CBF) is at a basal level. A constant oxygen extraction rate fueling neural activity leads to a fixed deoxygenated hemoglobin (Hb) to oxygenated hemoglobin (HbO₂) ratio in the capillary bed and venules. Since Hb is paramagnetic, it distorts the magnetic field. The Hb-related magnetic field inhomogeneities lead to rapid dephasing of excited spins

resulting in a low MRI signal level (*lower part*). **(b)** If the cortical region is in activated state, synaptic signal integration and spiking activity increases, leading to an increased oxygen extraction rate (*upper part*). CBF strongly increases delivering oxygen beyond local need, which essentially flushes Hb away from the capillary bed (*middle part*). Since HbO₂ does not substantially distort the homogeneity of the local magnetic field, excited spins dephase slower than in the baseline state (*lower part*) resulting in an enhanced MRI signal (BOLD effect)

all positions within a 2D image can be uniquely encoded with a desired resolution. Unfortunately a single application of the phase-encoding gradient is not sufficient to encode the second image dimension. The process of excitation and phase encoding must be repeated many times for a single slice. At each repetition, the strength of the phase-encoding gradient is slightly changed in order to ultimately obtain a complete frequency x phase encoding of the slice.

1.2.4 Two-Dimensional k Space

The data obtained from a series of excitation – recording cycles – can be arranged in a two-dimensional space called k space. Each row of k space corresponds to the data of one excitation – recording cycle with a different phase-encoding step. As described above, the echo signal of one line in k space contains a frequency-encoded representation of one dimension of the selected slice. While the slice selection and frequency-encoding gradients are the same from cycle to cycle, the slope of the phase-encoding gradient is changed by

a constant value across cycles and, thus, from line to line in k space. The imposed phase shift for a specific proton depends on the strength of the phase-encoding gradient and on the proton's position along the second image dimension. A series of phase-encoding steps “fills” k space in such a way that the second slice dimension ultimately also gets *frequency* encoded. The k space thus contains two-dimensional frequency-encoded information of the slice, which can be transformed into two-dimensional *image space* by application of the two-dimensional Fourier transform (2D FT).

1.2.5 Echo-Planar Imaging

The described procedure is applied for each slice of a scanned volume. A properly specified series of electromagnetic pulses allowing to construct one or more 2D images from electromagnetic echoes is called an *MRI pulse sequence*. The most often-used sequence for functional MRI is *gradient-echo echo-planar imaging* (GE-EPI). This sequence enables very rapid imaging of a slice by performing all phase-encoding

steps after a single 90° excitation pulse. This sequence requires switching the readout gradient rapidly on and off to fill k space line by line resulting in a series of (e.g., 64) small *gradient echoes* within the duration of a single T_2^* decay. A complete image can thus be obtained in about 50–100 ms as opposed to several seconds with standard (functional) imaging sequences. GE-EPI is very sensitive to field inhomogeneities influencing the speed of dephasing (T_2^* contrast). This is essential for functional imaging (see below) but also produces image distortions called susceptibility artifacts, which occur especially at tissue boundaries. Running EPI sequences requires a high-performance (i.e., expensive) gradient system to enable very rapid gradient switching.

1.2.6 Parallel Imaging and Parallel Excitation

In the last 15 years, *parallel imaging* (e.g., Pruessmann et al. 1999) has become a standard technique that has been introduced with different names by scanner manufacturers such as “SENSE,” “IPAT,” or “SMASH.” The basic idea of parallel imaging is the simultaneous acquisition of MRI data with at least two (typically 32 or more) receiver coils, each having a different spatial sensitivity. During image reconstruction, complementary information from the different receiver coils can be combined to fill k space in parallel reducing the number of time-consuming phase-encoding steps. Besides appropriate coils (phased array coils), parallel imaging requires that MRI scanners are equipped with multiple processing channels operating in parallel. Note that parallel imaging may be used either to increase temporal resolution when using a standard matrix size or to increase spatial resolution using a larger matrix with a conventional image acquisition time. Using parallel imaging to reduce scan time without sacrificing image quality is especially relevant for patient scans. Furthermore, parallel imaging may also reduce GE-EPI imaging artifacts because it allows acquiring standard image matrices with shorter echo times; typical EPI artifacts, such as signal dropouts in regions of neighboring tissue types and geometrical distortions, increase with increasing echo times.

In recent years, *parallel excitation* techniques are gaining increasing interest that work by exciting more than one slice in parallel: If, for example, eight slices are excited simultaneously, a whole-brain scan with 64 slices would be completed in the same time as eight nonsimultaneously recorded slices. In order to enable such powerful “multiband” techniques, an advanced excitation hardware (multiple transmit channels) is needed that is not yet standard on most MRI scanners. Furthermore, special MRI pulse sequences are needed (Moeller et al. 2010; Setsompop et al. 2012). Since multiple slices are acquired truly in parallel, imaging time is substantially reduced as compared to standard single-slice excitation techniques. This is especially beneficial for real-time fMRI neurofeedback studies (e.g., Goebel et al. 2010) since more time points (albeit temporally correlated)

can help to calculate more stable feedback values in a given time window. Note, however, that the data received from multiple slices need to be separated which becomes increasingly difficult with an increasing number of simultaneously excited slices. In order to avoid loss in image quality, the multiband factor (number of simultaneously excited slices) used for neuroscience applications is currently rather low, that is, in the range of 2–4.

2 Physiological Principles of fMRI

Neuronal activity consumes energy, which is produced by chemical processes requiring glucose and oxygen. The vascular system supplies these substances by a complex network of large and small vessels. The arterial part of the vascular system transports oxygenated blood through an increasingly fine-grained network of blood vessels until it reaches the capillary bed where the chemically stored energy (oxygen) is transferred to the neurons. If the brain is in resting state, 30–40 % of the oxygen is extracted from the blood in the capillary bed. The venous system transports the less-oxygenated blood away from the capillary bed. Oxygen is transported in the blood via the hemoglobin molecule. If hemoglobin carries oxygen, it is called oxygenated hemoglobin (HbO_2), while it is called deoxygenated hemoglobin (Hb) when it is devoid of oxygen. While the arterial network contains almost only oxygenated hemoglobin, the capillary bed and the venous network contain a mixture of oxygenated and deoxygenated hemoglobin.

2.1 Neurovascular Coupling

A local increase of neuronal activity immediately leads to an increased oxygen extraction rate in the capillary bed and, thus, in an increase in the relative concentration of deoxygenated hemoglobin. This fast response to increased neuronal activity is described as the “initial dip” (Fig. 7). After a short time of about 3 s, the increased local neuronal activity also leads to a strong increase in local blood flow. This response of the vascular system to the increased energy demand is called the *hemodynamic response*. Recent studies indicate that synaptic signal integration (measured by the local field potential, LFP) is a better predictor of the strength of the hemodynamic response than spiking activity (Logothetis et al. 2001; Mathiesen et al. 2000). It thus seems likely that the hemodynamic response primarily reflects the input and local processing of neuronal information rather than the output signals (Logothetis and Wandell 2004). Note that it is not yet completely known how the neurons “inform” the vascular system about their increased energy demand. Important theories about this *neurovascular coupling* are described, among many others, by Fox et al. (1988), Buxton

et al. (1998), and Magistretti et al. (1999). It appears likely that astrocytes play an important role because these special glial cells are massively connected with both neurons and the vascular system. The hemodynamic response consists in increased local cerebral blood flow (CBF) as well as increased cerebral blood volume (CBV), probably as a mechanical consequence of increased blood flow. The hemodynamic response not only compensates quickly for the slightly increased oxygen extraction rate but it is so strong that it results in a substantial local *oversupply* of oxygenated hemoglobin (Figs. 6 and 7). Note that it is not yet clear why the vascular system responds with a much stronger increase in cerebral blood flow than appears to be necessary. The increased CBV may help to explain the poststimulus undershoot (Fig. 7) observed in typical fMRI responses (balloon model, Buxton et al. 1998). While CBF and oxygen extraction rate may quickly return to baseline, the elastic properties of the dilated venules will require many seconds until baseline size is reached. In the expanded space of the dilated vessels, more deoxygenated hemoglo-

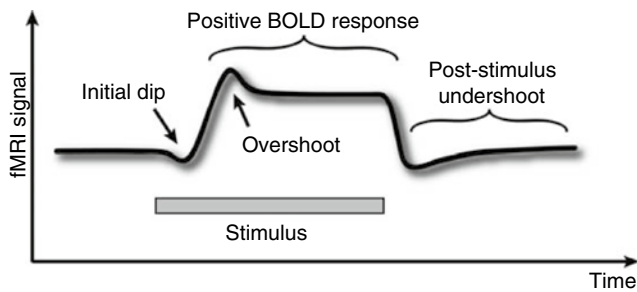


Fig. 7 Idealized time course of the hemodynamic response following a long (about 20 s) stimulation event. The theoretically expected initial dip is not reliably measured in human fMRI studies. For long stimulation events, the signal rises initially to a higher value (overshoot) than the subsequently reached plateau. When the stimulus is turned off, the signal often falls below the baseline signal level (undershoot), which is then approached slowly

bin will accumulate reducing the MRI signal below the pre-stimulus baseline level.


2.2 The BOLD Effect

The most common method of functional MRI is based on the BOLD effect (Ogawa et al. 1990). This exploits the fact that oxygenated hemoglobin has different magnetic properties than deoxygenated hemoglobin. More specifically, while oxygenated hemoglobin is diamagnetic, deoxygenated hemoglobin is paramagnetic altering the local magnetic susceptibility, creating magnetic field distortions within and around the blood vessels in the capillary bed and venules. During the hemodynamic response (oversupply phase), the oxygenated to deoxygenated hemoglobin ratio increases resulting in a more homogeneous local magnetic field. As follows from the description in Sect. 1, excited spins dephase slower in a more homogeneous magnetic field leading to a stronger measured MRI signal in the activated state when compared to a resting state (Fig. 6). The BOLD effect, thus, measures increased neuronal activity indirectly via a change in local magnetic field (in)homogeneity, which is caused by an oversupply of oxygenated blood (Fig. 6). Note that these field inhomogeneities are only detectable with MRI because of the different magnetic properties of oxy- and deoxygenated hemoglobin. The change in the local HbO_2/Hb ratio and its associated change in magnetic field homogeneity, thus, acts as an endogenous marker of neural activity.

2.3 The BOLD Hemodynamic Response

The time course of evoked fMRI signals, reflecting the BOLD hemodynamic response, is well studied for the primary visual cortex (V1). After application of a short visual stimulus of 100 ms, the observed (positive) signal response

Fig. 8 The two gamma function allows to model typical hemodynamic impulse responses. One gamma function models the peak (τ) and dispersion (σ) of the positive BOLD response, while the second gamma function models the peak and dispersion of the undershoot response. Parameter A scales the amplitudes of the individual gamma functions



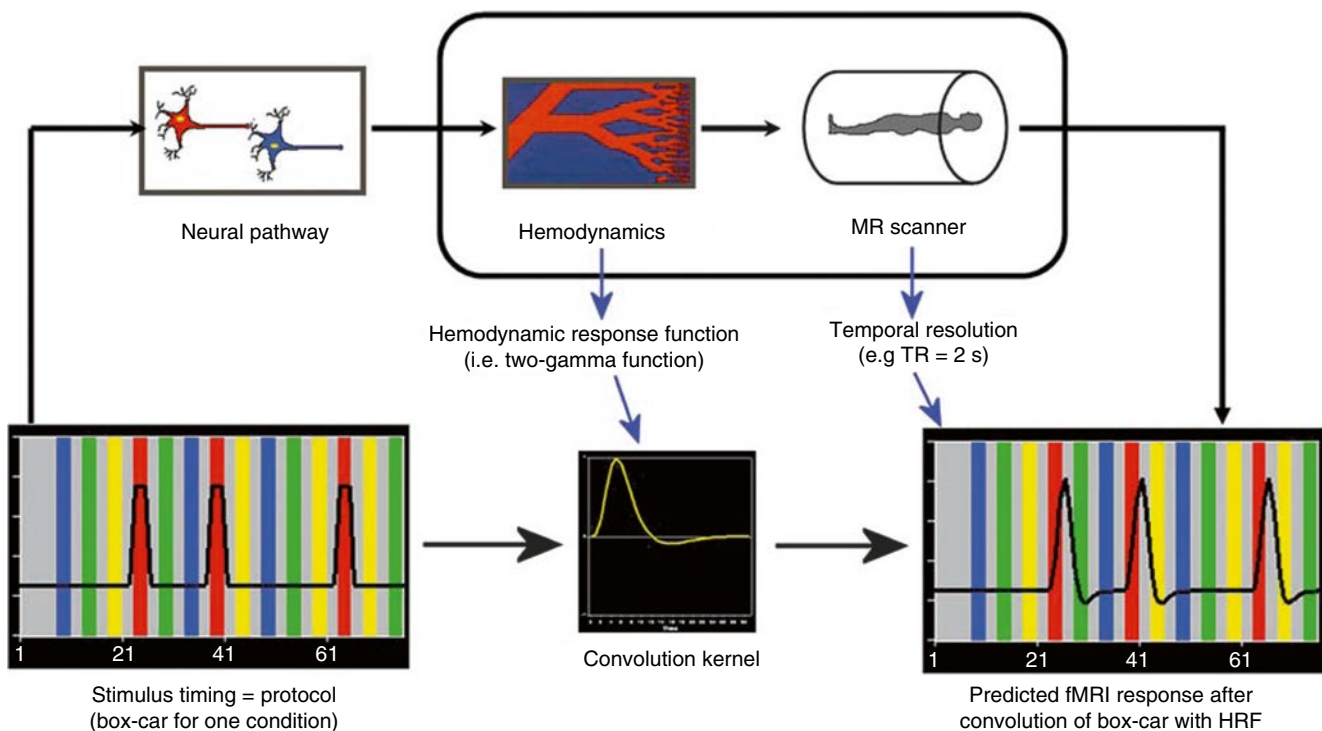


Fig. 9 Calculation of expected fMRI signal response for one condition of a protocol using convolution. The calculated response depends on the chosen model for the BOLD hemodynamic response function (HRF), for example, two gamma function (*middle part*). The expected

response is obtained by convolution of the box-car time course (*left*) with the chosen HRF. The convolved time course is downsampled to the temporal resolution (sampling intervals) of the fMRI measurements given by the volume TR value (*right*)

starts to rise after 2–3 s (oversupply phase) and reaches a maximum level after 5–6 s. About 10 s later, the signal reaches again the baseline level. As compared to the neuronal response of about 100 ms duration, the corresponding fMRI response is characterized by a delayed, gradual response profile extending as long as 20 s. Despite this sluggish response, the latency of response onsets appear to reflect quite precisely neuronal onset times (Menon and Kim 1999): If the left and right visual field are stimulated sequentially with a stimulus onset asynchrony of only 100 ms, response profiles from the right and left primary visual cortex are systematically shifted according to the applied temporal offset. More generally, the fMRI signal may reflect the flow of information processing across different brain areas as a sequence of shifted response profiles. Estimates of the temporal resolution with respect to onset delays are more in the order of hundreds of milliseconds than in the order of seconds (Formisano and Goebel 2003).

Assuming a linear time invariant (LTI) system, one can predict the expected time course of arbitrary long stimulation periods from the known response to a short stimulus. The response to a very short stimulus is called the *impulse response function* or, in the context of fMRI, the *BOLD hemodynamic response function* (HRF). The output (expected fMRI response) of an LTI system is the *convolution* of the input time course (e.g., stimulation “box-car”

time course) with the system’s response to an impulse function (Fig. 9). For primary visual cortex (V1), Boynton et al. (1996) showed that the measured responses to stimuli with varying amplitudes and durations could be indeed predicted well from the response profile obtained from a short visual stimulus. A well-suited function to model the hemodynamic impulse function is the probability density function (pdf) of the gamma distribution scaled by parameter A :

$$y(x; A, \tau, \sigma) = Ax^{\tau/\sigma-1} \frac{e^{-x/\sigma}}{\sigma^{\tau/\sigma} \Gamma(\tau/\sigma)}$$

Parameters τ and σ define the onset and dispersion of the response peak, respectively. While Boynton et al. (1996) used a single gamma function to characterize the impulse response function, the sum of two gamma functions (Friston et al. 1998) allows to also capture the undershoot usually observed in fMRI responses. The first gamma function typically peaks 5 s after stimulus onset ($\tau=6$), while the second gamma function peaks 15 s after stimulus onset ($\tau=16$, see Fig. 8). After convolution of a stimulus time course with the impulse function (Fig. 9), the calculated time course can be directly used as a reference function for statistical data analysis (see Sect. 3.3).

Note that the linear system assumption is reasonably valid only for stimuli of sufficiently long duration. For a series of

short stimuli separated by intervals shorter than 2–4 s, non-linear interaction effects have to be expected (e.g., Robson et al. 1998). Note further that the calculation (convolution) of expected time courses requires as input the valid specification of the time course of assumed neuronal response profiles, which is often not simply a copy of stimulus timing. A simple box-car time course, for example, assumes that neurons in a stimulated cortical area are active with constant amplitude in prolonged “on” periods. It is, however, well known that this assumption is too simplistic for neurons in early sensory areas. For higher cortical areas, for example, frontal areas involved in working memory, the neuronal response profile might differ substantially with respect to stimulus timing. Assuming that neuronal responses are correctly specified, it appears reasonable to use the same hemodynamic response function for all brain regions to predict expected BOLD signal time courses since neurovascular coupling should be similar in different brain areas. In case that it is difficult to specify proper input response profiles, a more general approach should be used (e.g., deconvolution analysis; see Sect. 3.3).

While fMRI responses clearly reflect the oversupply phase of the hemodynamic response, the theoretically expected initial dip (Fig. 7) has not been reliably detected in standard human fMRI measurements (for animal studies, see, e.g., Kim et al. 2000). This component of the idealized hemodynamic response is thus not included in the standard single or two gamma convolution kernels (Fig. 8). Data analysis of almost all fMRI studies is therefore based on the signals coming from the much stronger and sustained positive BOLD response.

2.4 Limits of Spatial and Temporal Resolution

The ultimate spatial and temporal resolution of fMRI is not primarily limited by technical constraints but by properties of the vascular system. The spatial resolution of the vascular system, and hence fMRI, seems to be in the order of 0.5–1 mm since relevant blood vessels run vertically through the cortex in roughly that distance (Duvernoy et al. 1981). An achievable resolution of 0.5–1 mm might be just enough to resolve *cortical columns*. A cortical column contains thousands of neurons possessing similar response specificity. A conventional brain area, such as the fusiform face area, could contain a set of cortical columns, each coding a different basic (e.g., face) feature. Cortical columns could, thus, form the basic building blocks (“alphabet”) of complex representations (Fujita et al. 1992). Since neurons within a column code for roughly the same feature, measuring the brain at the level of cortical columns promises to provide a relevant level for describing brain functioning. In cat visual cortex,

for example, orientation columns could be measured with fMRI at ultrahigh magnetic fields (4 and 9 T, Kim et al. 2000). The observed pattern of active orientation columns systematically changed when showing cats gratings of different orientations. Using ultrahigh magnetic fields (e.g., 7 T), columnar resolution appears to be within reach also for human brain imaging (e.g., Cheng et al. 2001; Yacoub et al. 2008; Zimmermann et al. 2011).

Despite the sluggishness of the fMRI signal, it has been shown that the obtained responses may reflect timing information with very high temporal precision. The signal of the left and right visual cortex, for example, reliably reflects temporal differences between stimulation of the left and right visual field as short as 100 ms (Menon and Kim 1999). When properly taking care of different hemodynamic delays in different brain areas, the analysis of BOLD onset latencies may also be very useful in revealing the sequential order of activity across brain areas within trials of complex cognitive tasks (fMRI mental chronometry, e.g., Formisano and Goebel 2003). In order to measure the brain with a temporal resolution in the order of milliseconds, other methods such as electroencephalography (EEG) and magnetoencephalography (MEG) must be used. If one succeeds in performing a proper combined analysis of EEG/MEG and fMRI data (Scherg et al. 1999; Dale and Halgren 2001; Bledowski et al. 2006), it becomes possible to describe brain function both with respect to its topographic distribution as well as with respect to its precise timing. While EEG/MEG data and fMRI data are conventionally obtained in different sessions, it has become possible to measure EEG data directly during fMRI recording sessions (e.g., Mulert et al. 2004).

3 fMRI Data Analysis

A major goal of functional MRI measurements is the localization of the neural correlates of sensory, motor, and cognitive processes. Another major goal of fMRI studies is the detailed characterization of the response profile for known regions-of-interest (ROIs) across experimental conditions. In this context, the aim of conducted studies is often not to map new functional brain regions (whole-brain analysis) but to characterize further how known specialized brain areas respond to (subtle) differences in experimental conditions (ROI-based analysis). Furthermore, it is often of interest to estimate the shape of the response and how it varies across different conditions and brain areas. Inspection of the shape of (averaged) time courses may also help to separate signal fluctuations due to measurement artifacts from stimulus-related hemodynamic responses. In order to obtain fMRI data with relatively high temporal resolution, functional time series are acquired using fast MR sequences sensitive to BOLD contrast. As described above most fMRI experiments

use the gradient echo EPI sequence, which allows acquisition of a 64×64 matrix in 50–100 ms. A typical functional scan of the whole brain with 20–40 slices lasts only 1–2 s on state-of-the-art MRI scanners. The data obtained from scanning all slices once at different positions (e.g., 30 slices covering the whole brain) is subsequently referred to as a *functional volume* or a *functional 3D image*. The measurement of an uninterrupted series of functional volumes is referred to as a *run*. A run, thus, consists of the repeated measurement of a functional volume and, hence, the repeated measurement of the individual slices. The sampling interval – the time until the same brain region is measured again – is called *volume TR*. The volume TR specifies the temporal resolution of the functional measurements since all slices comprising one functional volume are obtained once during that time. Note, however, that the slices of a functional volume are not recorded simultaneously, which implies that data from different regions of the brain are recorded at different moments in time (see “Slice Scan Time Correction” in Sect. 3.2). During a functional experiment, a subject performs tasks typically involving several experimental conditions. A short experiment can be completed in a single run, which typically consists of 100–1,000 functional volumes. Assuming a run with 500 volumes each consisting of 30 slices of 64×64 pixels and that two bytes are needed to store each pixel, the amount of raw data acquired per run would be $500 \times 30 \times 64 \times 64 \times 2 = 122,880,000$ bytes or roughly 117 MB. In more complex experiments, a subject typically performs multiple runs in one scanning session resulting in about 500 MB of functional data per subject per session. Using fast parallel imaging techniques and/or high-resolution scanning (e.g., slices with 128×128 pixels) several gigabytes (GBs) of raw image space data will be recorded per subject.

Given the small amplitude of task-related BOLD signal changes of typically 1–5 % and the presence of many confounding effects, such as signal drifts and head motion, the localization and characterization of brain regions responding to experimental conditions of the stimulation protocol is a nontrivial task. The major analysis steps of functional and associated anatomical data will be described in the following paragraphs including spatial and temporal *preprocessing*, *statistical data analysis*, *coregistration* of functional and anatomical data sets, and *spatial normalization*. Although these essential data analysis steps are performed in a rather standardized way in all major software packages, including AFNI (<http://afni.nimh.nih.gov/afni/>), BrainVoyager (<http://www.brainvoyager.com/>), FSL (<http://www.fmrib.ox.ac.uk/fsl/>), and SPM (<http://www.fil.ion.ucl.ac.uk/spm/>), there is still room for improvements as will be discussed below. For the visualization of functional data, high-resolution anatomical data sets with a resolution of (or close to) 1 mm in all three dimensions are often collected in a recording session.

In most cases, these anatomical volumes are scanned using slow $T1$ -weighted MR sequences that are optimized to produce high-quality images with very good contrast between the gray and white matter. In some analysis packages, anatomical data sets do not only serve as a structural reference for the visualization of functional information but are often also used to improve the functional analysis itself, for example, by restricting statistical data analysis to gray matter voxels or to analyze topological representations on extracted cortex meshes. The preprocessing of high-resolution anatomical data sets and their role in functional data analysis will be described in Sect. 3.4. Since some data analysis steps depend on the details of the experimental paradigm, the next section shortly describes the two most frequently used experimental designs.

3.1 Block- and Event-Related Designs

In the first years of fMRI measurements, experimental designs were adapted from positron-emission tomography (PET) studies. In the typical PET design, several *trials* (individual stimuli, or more generally, cognitive events) were clustered in *blocks*, each of which contained trials of the same condition (Fig. 10). As an example, one block may consist of a series of different pictures showing happy faces and another block may consist of pictures showing sad faces. The statistical analysis of such *block designs* compares the mean activity obtained in the different experimental blocks. Block designs were necessary in PET studies because of the limited temporal resolution of this imaging technique requiring about a minute to obtain a single whole-brain functional image. Since the temporal resolution of fMRI is much higher than PET, it has been proposed to use *event-related designs* (Blamire et al. 1992; Buckner et al. 1996; Dale and Buckner 1997). The characteristics of these designs (Fig. 10) follow closely those used in event-related potential (ERP) studies. In event-related designs, individual trials of different conditions are not clustered in blocks but are presented in a random sequence with sufficient time between trials to separate successive responses. Responses to trials belonging to the same condition are selectively averaged, and the calculated mean responses are statistically compared with each other. While block designs are well suited for many experiments, event-related designs offer several advantages over block designs, especially for cognitive tasks. An important advantage of event-related designs is the possibility to present stimuli in a randomized order (Fig. 10) avoiding cognitive adaptation or expectation strategies of the subjects. Such cognitive adaptations are likely to occur in block designs since a subject knows what type of stimuli to expect within a block after having experienced the first few trials. Another important advantage of event-related designs is that the *response profile* for different trial types




Fig. 10 In a block design (*upper row*), trials (events) belonging to the same condition are grouped together and are separated by a baseline block. In this example, two blocks of two main conditions (*green* – condition 1, *violet* – condition 2) are depicted. In slow event-related

designs, trials of different conditions appear in randomized order and are spaced sufficiently far apart to avoid largely overlapping BOLD responses. Optimal intertrial intervals (ITIs) are about 12 s

(and even single trials) can be *estimated* by event-related averaging. Furthermore, event-related designs allow *post hoc sorting* of individual brain responses. One important example of post hoc sorting is the separation of brain responses for correctly vs. incorrectly performed trials.

The possibilities of event-related fMRI designs are comparable to standard behavioral and ERP analyses. Note, however, that the hemodynamic response extends over about 20–30 s (Fig. 8) after presentation of a short stimulus; if only the positive BOLD response is considered, the signal extends over 10–15 s. The easiest way to conduct event-related fMRI designs is to temporally separate individual trials far enough to avoid overlapping responses of successive trials. Event-related designs with long temporal intervals between individual trials are termed *slow event-related designs* (Fig. 10). For stimuli of duration of 1–2 s, the optimal *intertrial interval* (ITI) for statistical analysis is about 12 s (Bandettini and Cox 2000; Maus et al. 2010a). Since it has been shown that the fMRI signals of closely spaced trials add up approximately linearly (Boynton et al. 1996; Dale and Buckner 1997, see Sect. 2.3), it is also possible to run experiments with inter trial intervals of 2–6 s. Designs with short temporal intervals between trials are called *rapid event-related designs* (Fig. 10). While the measured response of rapid event-related designs will contain a combination of overlapping responses from closely spaced trials, condition-specific event-related time courses can be isolated using deconvolution analysis. Deconvolution analysis works correctly only under the assumption of a linear system (see Sect. 3.2) and requires randomized intertrial intervals (“jitter”), which can be easily obtained by adding “null” (baseline) trials when trial sequences are created for an experiment. Note, however, that single-trial analyses are only possible when using a slow event-related

design. While adding null trials and simple permutations of trial types produce already good event sequences for rapid event-related designs, statistical power can be maximized by using more advanced randomization procedures (Wager and Nichols 2003; Maus et al. 2010b). In general, block- and event-related designs can be statistically analyzed using the same mathematical principles (see Sect. 3.3.3).

It is important to note that conventional fMRI data does not provide an absolute signal of brain activity limiting the quantitative interpretation of results. The major part of the signal amplitude is related to proton density and T_2 tissue contrast varying across brain regions within and between subjects. Small BOLD-related signal fluctuations, thus, neither have a defined origin nor a unit. In light of these considerations, signal strengths in main experimental conditions cannot be interpreted absolutely but have to be assessed relative to the signal strength in other main or control conditions within voxels. As a general control condition, many fMRI experiments contain a baseline (“rest,” “fixation”) condition with “no task” for the subject. Such simple control conditions allow analyzing brain activity that is common in multiple main conditions that would not be detectable when only comparisons between main conditions could be performed. More complex experimental (control) conditions differ from the main condition(s) only in specific cognitive component allowing isolating brain responses specific to that component.

Responses to main conditions are often expressed as percent signal change relative to a baseline condition. Furthermore, it is recommended to vary conditions within subjects – and even within runs – since the lack of an absolute signal level increases variability when comparing effects across runs, sessions, or subjects. Some experiments require a *between subjects* design, including comparisons of




Fig. 11 During functional MRI measurements, a set of slices, often covering the whole brain, is scanned repeatedly over time. Although the repeated slice measurements look almost identical, small task-related signal fluctuations may occur at different brain regions at different moments in time (a). To visualize these subtle fluctuations, the time course of any desired brain region (region-of-interest, ROI) may be depicted (b). The smallest separate brain region one can select to display a time course in a two-dimensional image (slice) is called pixel (picture element) while the smallest region in a three-dimensional “image” is called voxel (volume element)

responses between different subject groups, for example, males vs. females or treatment group vs. control group. Note that the BOLD signal measured with conventional fMRI may be affected by medications that modify the neurovascular coupling, for example, by increasing or decreasing baseline cerebral blood flow (CBF). In order to obtain more quantitative evaluation of activation responses, it is, thus, recommended for patient studies to combine standard BOLD measurements with CBF measurements using arterial spin labeling (ASL) techniques (e.g., Buxton et al. 2004).

3.2 Basic Analysis Steps

3.2.1 Two Views on fMRI Data Sets

In order to better understand different fMRI data analysis steps, two different views on the recorded four-dimensional data sets are helpful. In one view (Fig. 11a), the 4D data is

conceptualized as a *sequence of functional volumes* (3D images). This view is very useful to understand spatial analysis steps. During 3D motion correction, for example, each functional volume of a run is aligned to a selected reference volume by adjusting rotation and translation parameters. The second view focuses on time courses of individual voxels (“voxel” = “volume element” analogous to “pixel” = picture element). This second view (Fig. 11b) helps to understand those preprocessing and statistical procedures, which process *time courses of individual voxels*. Most standard statistical analysis procedures including the general linear model (GLM) operate in this way. In a GLM analysis, for example, the data is processed “voxel-wise” (univariate) by fitting a model to the time course of each voxel independently.

3.2.2 Preprocessing of Functional Data

In order to reduce artifact and noise-related signal components, a series of *preprocessing* operations is typically performed prior to statistical data analysis. The most essential *preprocessing* steps are (1) head motion detection and correction, (2) slice scan timing correction, (3) removal of linear and nonlinear trends in voxel time courses, and (4) spatial and temporal smoothing of the data.

3.2.2.1 Detection and Correction of Head Motion

The quality of fMRI data is strongly hampered in the presence of substantial head movements. Data sets are usually rejected for further analysis if head motion exceeds 5 mm. Although head motion can be corrected in image space, displacements of the head reduce the homogeneity of the magnetic field, which is fine-tuned (“shimmed”) prior to functional scans for the head position at that time. If head movements are small, 3D motion correction is an important step to improve data quality for subsequent statistical data analysis. Motion correction operates by selecting a functional volume of a run (or a volume from another run of the same scanning session) as a reference to which all other functional volumes are aligned. Most head motion algorithms describe head movements by six parameters assessing translation (displacement) and rotation at each time point with respect to the reference volume. These six parameters are appropriate to characterize motion of *rigid bodies*, since any spatial displacement of rigid bodies can be described by translation along the x -, y -, and z -axes and rotation around these axes. The values of these six parameters are estimated iteratively by analyzing how a source volume should be translated and rotated in order to better align with the reference volume; after applying a first estimate of the parameters, the procedure is repeated to improve the “fit” between the transformed (motion-corrected) and target (reference) volume. A *similarity* or *error measure* quantifies how good the transformed volume fits the reference volume. An often-used error measure is the sum of squared intensity differences at corresponding positions in the reference volume and the transformed

volume. The iterative adjustment of the parameter estimates stops if no further improvement can be achieved, that is, when the error measure reaches a minimum. After the final motion parameters have been *detected* by the iterative procedure, they can be applied to the source volume to produce a *motion-corrected* volume replacing the original volume in the output (motion-corrected) data set. For visual inspection, fMRI software packages are usually presenting line plots of the three translation and three rotation parameters across time showing how the estimated values change from volume to volume. The obtained parameter time courses may also be integrated in subsequent statistical data analysis with the aim to remove residual motion artifacts (for details, see Sect. 3.3).

Note that the assumption of a rigid body is not strictly valid for fMRI data since the individual slices of a functional volume are not scanned in parallel. Since abrupt head motions may occur at any moment in time, the assumption of a rigid body is violated. Imagine, for example, that a subject does not move while the first five slices of a functional volume are scanned, then moves 2 mm along the *y*-axis, and then lies still until scanning of that volume has been completed. The six parameters of a rigid body approach are not sufficient to capture such “within-volume” motion correctly. Fortunately, head movements from volume to volume are typically small and the assumption of a moving rigid body is, thus, largely valid.

3.2.2.2 Slice Scan Time Correction

For statistical analysis, a functional volume is usually considered as measured at the same time point. Individual slices (or a few slices when using state-of-the-art “multiband” sequences) of a functional volume are, however, scanned sequentially in standard 2D functional (EPI) measurements, that is, each slice (or set of slices in multiband sequences) is obtained at a different time point within a functional volume measurement. For a functional volume of 30 slices and a volume TR of 3 s, for example, the data of the last slice is measured almost 3 s later than the data of the first slice. Despite the sluggishness of the hemodynamic response (Fig. 8), an imprecise specification of time in the order of 3 s will lead to suboptimal statistical analysis, especially in event-related designs. It is, thus, desirable to preprocess the data in such a way that the resulting processed data appears as if all slices of a functional volume were measured at the same moment in time. Only then would it be, for example, possible to compare and integrate event-related responses from different brain regions correctly with respect to temporal parameters such as onset latency. In order to correct for different slice scan timings, the time series of individual slices are temporally “shifted” to match a reference time point, for example, the first or middle slice of a functional volume. The appropriate temporal shift of the time courses of the other slices is then performed by resampling the original data accordingly. Since this process involves sampling at time points that fall between measurement time points, the new values need to be estimated by interpolation of values from past and future time points (Fig. 12). The most

often-used interpolation methods are linear, cubic spline, and sinc interpolation. Note that the time points of slice scanning depend also on the *acquisition order* specified at the scanner console. Besides an ascending or descending order, slices are often scanned in an interleaved mode, that is, the odd slice numbers are recorded first followed by the even slice numbers. After appropriate temporal resampling, all slices within a functional volume of the new data set represent the same time point (Fig. 12) and can, thus, be statistically analyzed with the same hemodynamic response function; if slice scan time correction is not performed, hemodynamic response functions should be adjusted (shifted) on a per-slice basis.

3.2.2.3 Removal of Drifts and Temporal Smoothing of Voxel Time Series

Due to physical and physiological noise, voxel time courses are often nonstationary exhibiting signal drifts over time. If the signal rises or falls with a constant slope from beginning to end of a run, the drift is described as a *linear trend*. If the signal level slowly varies over time with a nonconstant slope, the drift is described a nonlinear trend. Since drifts describe slow signal changes, they can be removed by Fourier analysis using a *temporal high-pass filter*. The original signal in the time domain is transformed in frequency space using the Fourier transform (FT). In the frequency-domain drifts can be easily removed because low-frequency components, underlying drifts, are isolated from higher-frequency components reflecting task-related signal changes. After applying a high-pass filter in the frequency domain (removing low frequencies), the data is transformed back into the time domain by the inverse Fourier transform (Fig. 13). As an alternative approach, drifts can be modeled and removed in the time domain using appropriate basis sets in a general linear model analysis. This approach can be performed either as a preprocessing step or as part of statistical data analysis (for details, see Sect. 3.3.3). Removal of drifts is recommended as a preprocessing step since it is not only relevant for statistical data analysis but also for the calculation of event-related time courses.

While less important, another temporal preprocessing step is *temporal smoothing* of voxel time courses removing *high-frequency signal fluctuations*, which are considered as noise. While this step increases the signal-to-noise ratio, temporal smoothing is not recommended when analyzing event-related designs since it may distort estimates of temporally relevant parameters, such as the onset or width of average event-related responses. Temporal smoothing also increases serial correlations between values of successive time points that need to be corrected (see Sect. 3.3.3.6).

3.2.2.4 Spatial Smoothing

To further enhance the signal-to-noise ratio, the data is often spatially smoothed by convolution with a 3D Gaussian kernel. In this process, each voxel is replaced by a weighted




Fig. 12 During slice scan time correction, slices within each functional volume (*black rectangles*) are “shifted in time” resulting in a new time series (*violet rectangles*) in which all slices of a functional volume are virtually measured at the same moment in time. To calculate intensity values at time points falling in-between measured time points, past and

future values have to be integrated typically using sinc or linear interpolation. For correct interpolation, the volume TR, slice TR, and slice scanning order must be known. (a) Five slices are scanned in ascending order. (b) Five slices are scanned in interleaved order

average value calculated across neighboring voxels. The shape and width of the Gaussian kernel determines the weights used to multiply the values of voxels in the neighborhood, that is, weights decrease with increasing distance from the considered voxel; voxels further apart from the center will, thus, contribute less to the weighted average than voxels close to the center of the considered voxel. Note that smoothing reduces the spatial resolution of the data and should be therefore applied with care. Many studies, however, aim to detect regions larger than a few voxels, that is, brain areas in the order of 1 cm^3 or larger. Under these conditions, spatial smoothing with an appropriate kernel width of 4–8 mm is useful since it suppresses noise and enhances task-related signals. Spatial smoothing also increases the extent of activated brain regions, which is exploited in the context of group analyses (see Sect. 3.5) facilitating the integration of signals from corresponding but not perfectly aligned brain regions.

From the description and discussion of standard preprocessing steps, it should have become clear that there are no universally correct criteria to choose preprocessing steps and parameters because choices depend to some extent on the goal of data analysis. Some steps depend also on the experimental design of a study. If, for example, a high-pass temporal filter is used with a cutoff point that is too high,

interesting task-related signal fluctuations could easily be removed accidentally from the data.

Besides the described core preprocessing steps, additional procedures may be applied. The next sections will describe three additional preprocessing steps.

3.2.2.5 Mean Intensity Adjustment

Besides drifts in *individual* voxel time courses, the mean intensity level averaged across all voxels might exhibit drifts over time. These global drifts can be corrected by scaling the intensity values of a functional volume in such a way that the new mean value is identical to the mean intensity value of a reference volume. Mean intensity adjustment is not strictly necessary since modern scanners keep a rather constant mean signal level over time. Under this condition, mean intensity adjustment may even produce a negative effect by reducing true activation effects. If, for example, large parts of the brain activate during a main condition as compared to a rest condition, the mean signal level is higher during active periods, and a mean intensity adjustment step will “correct” this. A plot of the mean signal level over time might be, however, helpful to identify problems of the scanner quality, especially when such a plot shows “spikes,” that is, strong signal decreases (or increases) at isolated time points.

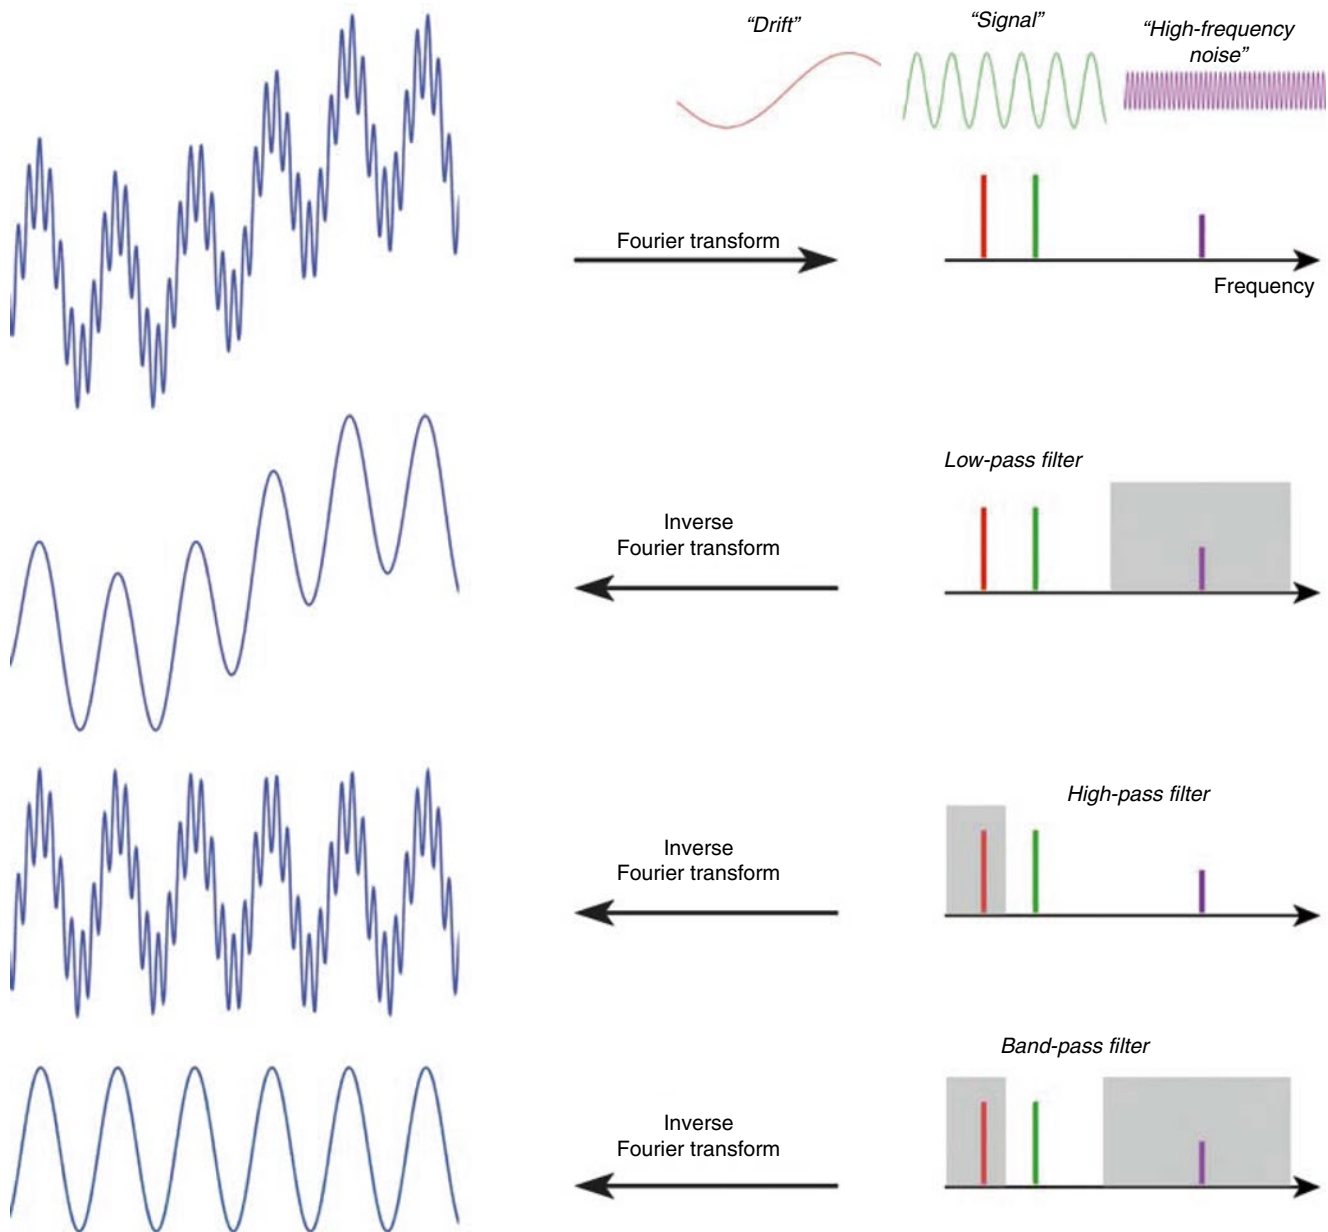


Fig. 13 Principle of temporal filtering using Fourier analysis. The time-domain signal can be converted in an equivalent frequency-domain signal using the Fourier transform (*upper row*). In this simplified example, the composite signal (*upper row, left*) consists only of three frequencies representing a drift, signal, and high-frequency noise

component (*upper row, right*). In the frequency domain, frequencies can be filtered to remove unwanted signal components. The filtered signal can then be converted back into the time domain using the inverse Fourier transform. In the second row, a low-pass filter is applied, in the third row a high-pass filter, and in the fourth row a band-pass filter

3.2.2.6 Motion Correction Within and Across Runs

A scanning session typically consists of a series of runs. In such a situation, head movements may not only occur within runs but also between runs. A simple approach to align all functional volumes of all runs of a scanning session with each other consists in specifying the same reference volume for all runs. If a session consists, for example, of three runs, all functional volumes could be aligned to the middle volume of the

second run. Since functional data is often aligned with a 3D anatomical data set recorded in the same session, it is recommended to choose a functional volume as a reference, which is recorded just before (or after) the anatomical data set. Note, however, that across-run motion correction works only if the slice positions are specified identically in all runs. If across-run motion correction is not possible, each run can also be individually adjusted to a common 3D anatomical data set.

3.2.2.7 Distortion Correction of Functional Images

The BOLD sensitive GE-EPI sequence is used for most fMRI studies because of its speed, but it has the disadvantage that images suffer from signal dropouts and geometric distortions, especially in brain regions close to other tissue types such as air and liquor (susceptibility artifacts). These artifacts can be reduced substantially by using optimized EPI sequence parameters (e.g., Weiskopf et al. 2006) and parallel imaging techniques (see Sect. 1.3). A complete removal of dropouts and geometric distortions is, however, not possible. Further improvements may be obtained by distortion correction routines, which may benefit from special scans measuring magnetic field distortions (e.g., field maps). The distortion-corrected images may improve coregistration results between functional and anatomical data sets enabling a more precise localization of brain function.

3.3 Statistical Analysis of Functional Data

Statistical data analysis aims at identifying those brain regions exhibiting increased or decreased responses in specific experimental conditions as compared to other (e.g., control) conditions. Due to the presence of physiological and physical noise fluctuations, observed differences between conditions might occur simply by chance. Note that measurements provide only a *sample* of data, but we are interested in true effects in the underlying *population*. At the level of individual functional scans, time points are treated as subjects, that is, sample corresponds to the obtained repeated measurements at every TR and “population” refers to the estimated but unobservable true condition effects within the subject. In multi-subject (group) analyses, sample usually refers to estimated effects obtained in each subject and population refers to all people from which the sample of subjects has been drawn. Statistical data analysis protects from wrongly accepting effects in small sample data sets by explicitly assessing the effect of measurement variability (noise fluctuations) on estimated condition effects: If it is very unlikely that an observed effect is solely the result of noise fluctuations, it is assumed that the observed effect reflects a true difference between conditions in the population. In standard single-subject statistical fMRI analyses, this assessment is usually performed independently for the time course of each voxel (univariate analysis). The obtained statistical values, one for each voxel, form a three-dimensional *statistical map*. In more complex analyses, each voxel will contain several statistical values reflecting estimated effects of multiple conditions. Since independent testing at each voxel increases the chance to find some voxels with strong differences between conditions simply due to noise fluctuations, further adjustments for *multiple comparisons* need to be made.

3.3.1 From Image Subtraction to Statistical Comparison

Figure 14 shows two fMRI time courses obtained from two different brain areas of an experiment with two conditions, a control condition (“Rest”) and a main condition (“Stim”). Each condition has been measured several times.¹ How can we assess whether the response values are higher in the main condition than in the control condition? One approach consists in subtracting the mean value of the “Rest” condition, \bar{X}_1 , from the mean value of the “Stim” condition, \bar{X}_2 that is: $d = \bar{X}_2 - \bar{X}_1$. Note that in this example, one would obtain the same mean values in both conditions and, thus, the same difference in cases (a) and (b). Despite the fact that the means are identical in both cases, the difference in case (b) seems to be more “trustworthy” than the difference in case (a) because the measured values exhibit less fluctuations, that is, they *vary less* in case (b) than in case (a).

Statistical data analysis goes beyond simple subtraction by taking into account the amount of variability of the measured data points. Statistical analysis essentially asks how likely it is to obtain a certain effect (e.g., difference of condition means) in a data sample if there is no effect at the population level, that is, how likely it is that an observed sample effect is solely the result of noise fluctuations. This is formalized by the *null hypothesis* stating that there is no effect, for example, no true difference between conditions in the population. In the case of comparing the two means μ_1 and μ_2 , the null hypothesis can be formulated as $H_0: \mu_1 = \mu_2$. Assuming the null hypothesis, it can be calculated how likely it is that an observed sample effect would have occurred simply by chance. This requires knowledge about the amount of noise fluctuations (and its distribution), which can be estimated from the data. By incorporating the number of data points and the variability of measurements, statistical data analysis allows to *estimate the uncertainty of effects* (e.g., mean differences) in data samples. If an effect is large enough so that it is very unlikely that it has occurred simply by chance (e.g., the probability is less than $p=0.05$), one *rejects* the null hypothesis and *accepts* the *alternative hypothesis* stating that there exists a true effect in the population. Note that the decision to accept or reject the null hypothesis is based on a probability value which has been accepted by the scientific community ($p < 0.05$). Since the decision to accept or reject the null hypothesis is based on a probability value, a statistical analysis does not *prove* the existence of an effect, it only suggests “to believe in an effect” if it is very unlikely that the observed effect has

¹Note that in a real experiment, one would not just present once the control and main condition as in Fig. 14, but several “on-off” cycles; with too few repetitions, task-related response could not be distinguished from potential low-frequency drifts (see Sect. 3.2.2).

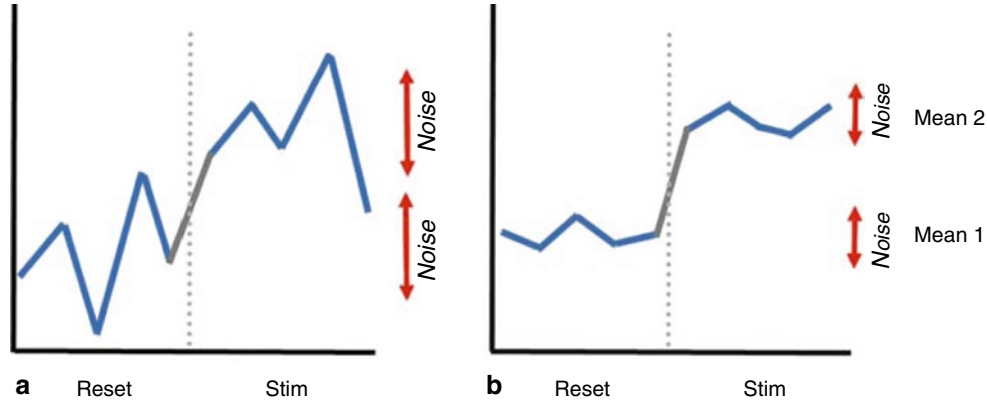


Fig. 14 Principle of statistical data analysis. An experiment with two conditions (“Stim” and “Rest”) has been performed. **(a)** Time course obtained in area 1. **(b)** Time course obtained in area 2. Calculation and subtraction of mean 1 (“Rest” condition) from mean 2 (“Stim” condition) leads to the same result in **(a)** and **(b)**. In a statistical analysis, the

estimated effect (mean difference) is related to its uncertainty, which is estimated by the variability of the measured values within conditions. Since the variance within the two conditions is smaller in **(b)** than in **(a)**, the estimated effect is more likely to correspond to a true difference in **(b)** than in **(a)**

occurred by chance. Note that a probability of $p=0.05$ means that if we would repeat the experiment 100 times, we would accept the alternative hypothesis in about five cases even if there would be no real effect in the population. Since the chosen probability value thus reflects the likelihood of wrongly rejecting the null hypothesis, it is also called *error probability*. The error probability is also referred to as the *significance level* and denoted with the Greek letter α . If one would know that there is no effect in the population but one would incorrectly reject the null hypothesis in a particular data sample, a “false-positive” decision would be made (type 1 error, “false alarm”). Since a false-positive error depends on the chosen error probability, it is also referred to as alpha error. If one would know that there is a true effect in the population but one would fail to reject the null hypothesis in a sample, a “false-negative” decision would be made, that is, one would miss a true effect (type 2 error).

3.3.2 t-Test and Correlation Analysis

The uncertainty of an effect is estimated by calculating the variance of the noise fluctuations from the data. For the case of comparing two mean values, the observed difference of the means is related to the variability of that difference resulting in a t statistic:

$$t = \frac{\bar{X}_2 - \bar{X}_1}{\hat{\sigma}_{\bar{x}_2 - \bar{x}_1}}$$

The numerator contains the calculated mean difference while the denominator contains the estimate of the expected variability, the *standard error* of the mean difference. Estimation of the standard error $\hat{\sigma}_{\bar{x}_2 - \bar{x}_1}$ involves pooling of the variances obtained within both conditions. Since we observe a high

variability in case (a) of the example data (Fig. 14), we will obtain a small t value. Due to the small variability of the data points in (b), we will obtain a larger t value in this case (Fig. 14). The higher the t value, the less likely it is that the observed mean difference is just the result of noise fluctuations. It is obvious that measurement of many data points allows a more robust estimation of this probability than the measurement of only a few data points. The error probability p can be calculated exactly from the obtained t value using the incomplete beta function $I_x(a,b)$ and the number of measured data points N :

$$p = I_{\frac{N-2}{N-2+t^2}}\left(\frac{N-2}{2}, \frac{1}{2}\right)$$

If the computed error probability falls below the standard value ($p < 0.05$), the alternative hypothesis is accepted stating that the observed mean difference exists in the population from which the data points have been drawn (i.e., measured). In that case, one also says that the two means differ *significantly*. Assuming that in our example the obtained p value falls below 0.05 in case (b) but not in case (a), we would only infer for brain area 2 that the “Stim” condition differs significantly from the “Rest” condition.

The described mean comparison method is not the ideal approach to compare responses between different conditions since this approach is unable to capture the gradual rise and fall of fMRI responses, for example, when a voxel exhibits a strong response to a trial of condition B after having not responded strongly to a preceding trial of condition A. As long as the temporal sampling resolution is low (volume TR > 4 s), the mean of different conditions can be calculated easily because transitions of expected responses from different conditions occur within a single time point (Fig. 15). If the

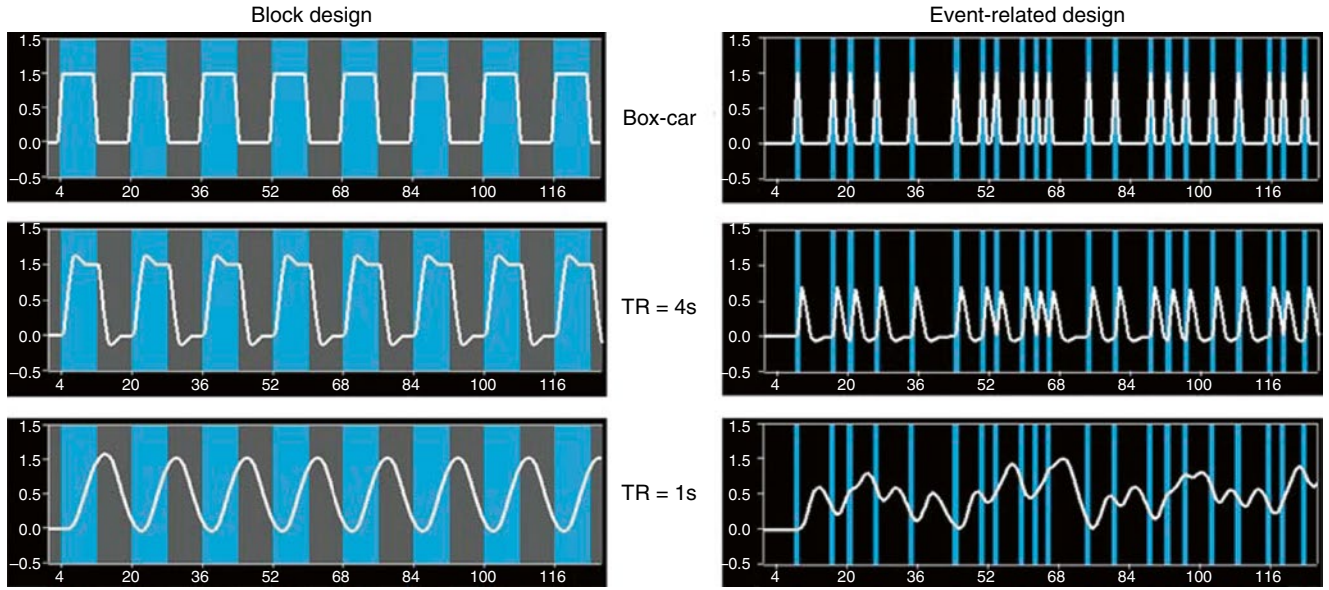


Fig. 15 Calculation of expected fMRI signals for a block- and event-related design. The horizontal axis of each plot represents time (data points). The vertical axis represents the amplitude of the modeled fMRI response. The blue vertical segments correspond to intervals of a single main stimulation condition; the gray segments correspond to a control condition. White curves show predicted BOLD responses. The plots in the upper row depict time courses, which do not take into account the delayed hemodynamic response

profile (“box-car”). The white curves in the other plots represent the expected time courses after application of a standard hemodynamic response function (two gamma function) for a temporal resolution (volume TR) of 4 s (middle row) and 1 s (lower row). Correlation analysis is able to capture the gradual increase and decrease of expected time courses for short TRs while it is impossible to unambiguously categorize time points as belonging to stimulation vs. baseline conditions in the context of a t test

temporal resolution is high, the expected fMRI responses change gradually from one condition to the next due to the sluggishness of the hemodynamic response (Fig. 15, TR = 1 s). In this case, time points in the “transitional zone” cannot be assigned easily to different conditions. Without special treatment, the mean response can no longer be easily computed for each condition. As a consequence, the statistical power to detect mean differences may be substantially reduced, especially for short blocks and events.

This problem does not occur when *correlation analysis* is used since this method allows explicitly incorporating the gradual increase and decrease of the expected BOLD signal. The predicted ideal (noise-free) time courses in Fig. 15 can be used as the *reference function* in a correlation analysis. At each voxel, the time course of the reference function is compared with the time course of the measured data from a voxel by calculating the *correlation coefficient* r , indicating the strength of covariation:

$$r = \frac{\sum_{t=1}^N (X_t - \bar{X})(Y_t - \bar{Y})}{\sqrt{\sum_{t=1}^N (X_t - \bar{X})^2 \sum_{t=1}^N (Y_t - \bar{Y})^2}}$$

Index t runs over time points (t for “time”) identifying pairs of temporally corresponding values from the reference (X_t)

and data (Y_t) time courses. In the numerator the mean of the reference and data time course is subtracted from the respective value of each data pair and the two differences are multiplied. The resulting value is the sum of these cross products, which will be high if the two time courses *covary*, that is, if the values of a pair are both above or below their respective means in most cases. The term in the denominator normalizes the covariation term in the numerator so that the correlation coefficient lies in a range of -1 and $+1$. A value of $+1$ indicates that the reference time course and the data time course go up and down in exactly the same way, while a value of -1 indicates that the two time courses run in opposite direction (anticorrelation). A correlation value of 0 indicates that the two time courses do not covary, that is, the value in one time course cannot be used to predict the corresponding value in the other time course.

While the statistical logic is the same in correlation analysis as described for mean comparisons, the null hypothesis now corresponds to the statement that the population correlation coefficient ρ equals zero ($H_0: \rho=0$). By including the number of data points N , the error probability can be computed assessing how likely it is that an observed correlation coefficient would occur solely due to noise fluctuations in the signal time course. If this probability falls below 0.05, the alternative hypothesis is accepted stating that there is indeed significant covariation between the reference function

and the data time course. Since the reference function is the result of a model assuming different response strengths in the two conditions (e.g., “Rest” and “Stim”), a significant correlation coefficient indicates that the two conditions lead indeed to different mean activation levels in the respective voxel or brain area. The statistical assessment can be performed also by converting an observed r value into a corresponding t value, $t = r\sqrt{N-2} / \sqrt{1-r^2}$.

3.3.3 The General Linear Model

The described t test for assessing the difference of two mean values is a special case of an analysis of a *qualitative (categorical) independent variable*. A qualitative variable is defined by discrete levels, for example, “stimulus on” vs. “stimulus off” or “male” vs. “female.” If a design contains more than two levels, a more general method such as analysis of variance (ANOVA) need to be used, which can be considered as an extension of the t test to more than two levels and to more than one experimental factor. The described correlation coefficient on the other hand is suited for the analysis of *quantitative independent variables*. A quantitative variable may be defined by any gradual time course. If more than one reference time course has to be considered, *multiple regression analysis* can be used, which can be considered as an extension of the simple linear correlation analysis. The *general linear model*² (GLM) is mathematically identical to a multiple regression analysis but stresses its suitability for both multiple qualitative and multiple quantitative variables. The GLM is suited to implement any parametric statistical test with one dependent variable, including any factorial ANOVA design as well as designs with a mixture of qualitative and quantitative variables (covariance analysis, ANCOVA). Because of its flexibility to incorporate multiple quantitative and qualitative independent variables, the GLM has become the core tool for fMRI data analysis after its introduction into the neuroimaging community by Friston and colleagues (Friston et al. 1994, 1995). The following sections briefly describe the mathematical background of the GLM in the context of fMRI data analysis; a comprehensive treatment of the GLM can be found in the standard statistical literature, for example, Draper and Smith (1998) and Kutner et al. (2005).

From the perspective of multiple regression analysis, the GLM aims to “explain” or “predict” the variation of a dependent variable in terms of a *linear combination* (weighted sum) of several reference functions. The dependent variable corresponds to the observed fMRI time course of a voxel and the reference functions correspond to time courses of expected (noise-free) fMRI responses for different conditions

of the experimental paradigm. The reference functions are also called *predictors*, *regressors*, *explanatory variables*, *covariates*, or *basis functions*. A set of specified predictors forms the *design matrix*, also called the *model*. A predictor time course is typically obtained by convolution of a “box-car” time course with a standard hemodynamic response function (Figs. 8 and 15). A box-car time course is usually defined by setting values to 1 at time points at which the modeled condition is defined (“on”) and 0 at all other time points.

Each predictor time course X_i gets an associated coefficient or *beta weight* b_i that quantifies the contribution of a predictor in explaining variance in the voxel time course y . The voxel time course y is modeled as the sum of the defined predictors, each multiplied with the associated beta weight b . Since this linear combination will not perfectly explain the data due to noise fluctuations, an error value e is added to the GLM system of equations with n data points and p predictors:

$$\begin{aligned} y_1 &= b_0 + b_1 X_{11} + \dots + b_p X_{1p} + e_1 \\ y_2 &= b_0 + b_1 X_{21} + \dots + b_p X_{2p} + e_2 \\ &\vdots \qquad \qquad \qquad \ddots \qquad \qquad \qquad \vdots \\ y_n &= b_0 + b_1 X_{n1} + \dots + b_p X_{np} + e_n \end{aligned}$$

The y variable on the left side corresponds to the data, that is, the measured time course of a single voxel. Time runs from top to bottom, that is, y_1 is the measured value at time point 1, y_2 the measured value at time point 2, and so on. The voxel time course (left column) is “explained” by the terms on the right side of the equation. The first column on the right side corresponds to the first beta weight b_0 . The corresponding predictor time course X_0 has a value of 1 for each time point and is, thus, also called “constant.” Since multiplication with 1 does not alter the value of b_0 , this predictor time course (X_0) does not explicitly appear in the equation. After estimation (see below), the value of b_0 typically represents the signal level of the baseline condition and is also called intercept. While its absolute value is not very informative in the context of fMRI data, it is important to include the constant predictor in a design matrix since it allows the other predictors to model small condition-related fluctuations as increases or decreases relative to the baseline signal level. The other predictors on the right-side model the expected time courses of different conditions. For multifactorial designs, predictors may be defined coding combinations of condition levels in order to estimate main and interaction effects. The beta weight of a predictor scales the associated predictor time course and reflects the unique contribution of that predictor in explaining part of the variance in the voxel time course. While the exact interpretation of beta values depends on the details of the design matrix, a large positive (negative) beta weight typically indicates that the voxel exhibits strong activation (deactivation) during the modeled experimental condition relative to baseline. All beta values together characterize a voxel’s “preference” for one or more experimental

²In the fMRI literature, the term “general linear model” refers to its univariate version where “univariate” refers to the number of dependent variables (one). In its general form, the general linear model has been defined for multiple dependent variables, that is, it encompasses tests as general as multivariate covariance analysis (MANCOVA).

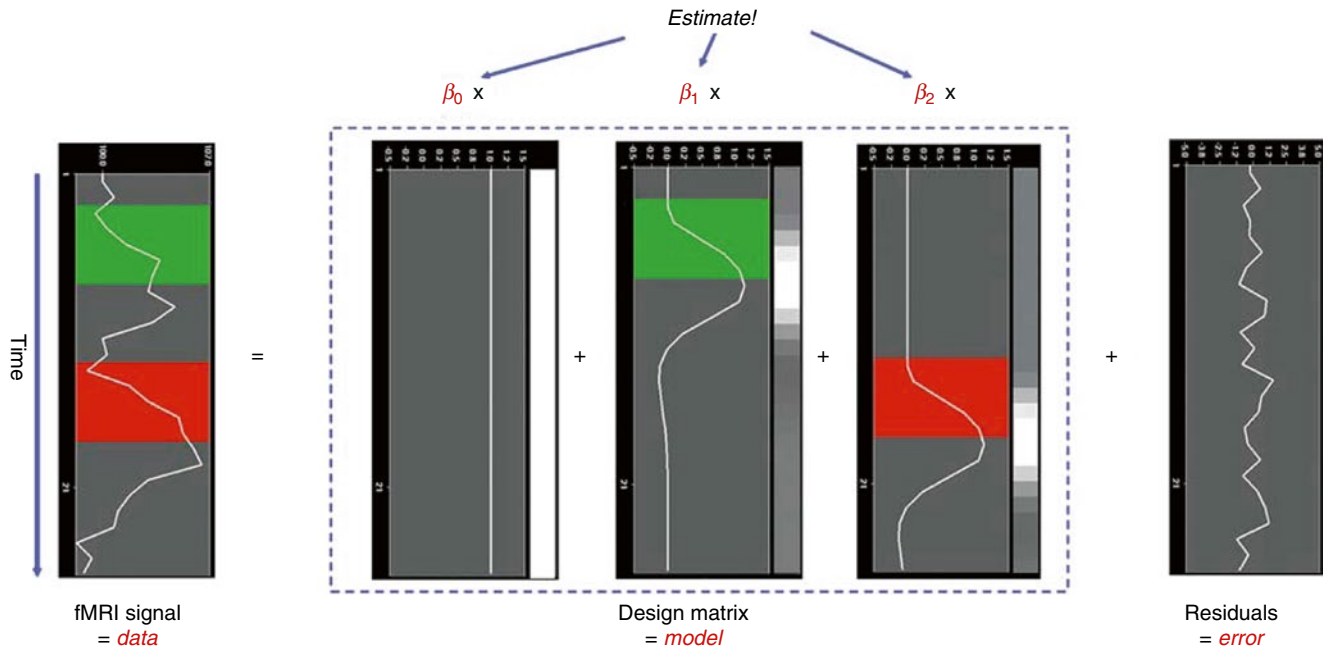


Fig. 16 Graphical display of a general linear model. Time is running from *top to bottom*. *Left side* shows observed voxel time course (data). The model (design matrix) consists of three predictors, the constant and two main predictors (*middle part*). Filled *green and red rectangles* depict stimulation time while the *white curves* depict expected BOLD responses. Expected responses are also shown in

graphical view using a *black-to-white* color range (*right side* of each predictor plot). Beta values have to be estimated (*top*) to scale the expected responses (predictors) in such a way that their weighted sum predicts the data values as good as possible (in the least squares sense, see text). Unexplained fluctuations (residuals, error) are shown on the *right side*

conditions. The last column in the system of equations contains error values, also called *residuals*, *prediction errors*, or *noise*. These error values quantify the deviation of the measured voxel time course from the predicted time course.

The GLM system of equations may be expressed elegantly using matrix notation. For this purpose, the voxel time course, the beta values, and the residuals are represented as vectors and the set of predictors as a matrix:

$$\begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} = \begin{bmatrix} 1X_{11} & \cdots & X_{1p} \\ \vdots & \vdots & \vdots \\ 1X_{n1} & \cdots & X_{np} \end{bmatrix} \begin{bmatrix} b_0 \\ b_1 \\ \vdots \\ b_p \end{bmatrix} + \begin{bmatrix} e_1 \\ e_2 \\ \vdots \\ e_n \end{bmatrix}$$

Representing the indicated vectors and matrix with single letters, we obtain this simple form of the GLM system of equations:

$$\mathbf{y} = \mathbf{Xb} + \mathbf{e}$$

In this notation, the matrix \mathbf{X} represents the design matrix containing the predictor time courses as column vectors. The beta values now appear in a separate vector \mathbf{b} . The term \mathbf{Xb} indicates matrix-vector multiplication. Figure 16 shows a graphical representation of the GLM. Time courses of the signal, predictors, and residuals have been arranged in column form with time running from top to bottom as in the system of equations.

Given the data \mathbf{y} and the design matrix \mathbf{X} , the GLM fitting procedure has to find a set of beta values explaining the data as good as possible. The time course values $\hat{\mathbf{y}}$ predicted by the model are obtained by the linear combination of the predictors:

$$\hat{\mathbf{y}} = \mathbf{Xb}$$

A good fit would be achieved with beta values leading to predicted values $\hat{\mathbf{y}}$ that are as close as possible to the measured values \mathbf{y} . By rearranging the system of equations, it is evident that a good prediction of the data implies small error values:

$$\begin{aligned} \mathbf{e} &= \mathbf{y} - \mathbf{Xb} \\ &= \mathbf{y} - \hat{\mathbf{y}} \end{aligned}$$

An intuitive idea would be to find those beta values minimizing the sum of error values. Since the error values contain both positive and negative values (and because of additional statistical considerations), the GLM procedure does not estimate beta values minimizing the sum of error values but finds those beta values *minimizing the sum of squared error values*:

$$\mathbf{e}'\mathbf{e} = (\mathbf{y} - \mathbf{Xb})'(\mathbf{y} - \mathbf{Xb}) \rightarrow \min$$

The term $\mathbf{e}'\mathbf{e}$ is the vector notation for the sum of squares $\left(\sum_{i=1}^N e_i^2 \right)$. The apostrophe symbol denotes transposition of a

vector (or matrix), that is, a row vector version of \mathbf{e} is multiplied by a column vector version of \mathbf{e} resulting in the sum of squared error values e_r . The optimal beta weights minimizing the squared error values (the “least squares estimates”) are obtained non-iteratively by the following equation:

$$\mathbf{b} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{y}$$

The term in brackets contains a matrix-matrix multiplication of the transposed design matrix \mathbf{X}' and the non-transposed design matrix \mathbf{X} . This term results in a square matrix with a number of rows and columns corresponding to the number of predictors. Each cell of the $\mathbf{X}\mathbf{X}$ matrix contains the scalar product of two predictor vectors. The scalar product is obtained by summing all products of corresponding entries of two vectors corresponding to the (non-mean normalized) calculation of covariance. This $\mathbf{X}\mathbf{X}$ matrix, thus, corresponds to the (non-mean normalized) predictor variance-covariance matrix.

The resulting square matrix is inverted as denoted by the “ -1 ” symbol. The resulting matrix $(\mathbf{X}\mathbf{X})^{-1}$ plays an essential role not only for the calculation of beta values but also for testing the significance of contrasts (see below). The remaining term on the right side, $\mathbf{X}'\mathbf{y}$, evaluates to a vector containing as many elements as predictors. Each element of this vector is the scalar product (non-mean normalized covariance term) of a predictor time course with the observed voxel time course.

An important property of the least squares estimation method (following from the independence assumption of the errors, see below) is that the variance of the measured time course can be decomposed into the sum of the variance of the predicted values (model-related variance) and the variance of the residuals:

$$\text{Var}(\mathbf{y}) = \text{Var}(\hat{\mathbf{y}}) + \text{Var}(\mathbf{e})$$

Since the variance of the voxel time course is fixed, minimizing the error variance by least squares corresponds to maximizing the variance of the values explained by the model. The square of the *multiple correlation coefficient* R provides a measure of the proportion of the variance of the data which can be explained by a specified model:

$$R^2 = \frac{\text{Var}(\hat{\mathbf{y}})}{\text{Var}(\mathbf{y})} = \frac{\text{Var}(\hat{\mathbf{y}})}{\text{Var}(\hat{\mathbf{y}}) + \text{Var}(\mathbf{e})}$$

The values of the multiple correlation coefficient vary between 0 (no variance explained) and 1 (all variance explained by the model). A coefficient of $R=0.7$, for example, corresponds to an explained variance of 49 % (0.7^2). An alternative way to calculate the multiple correlation coefficient consists in computing a standard correlation coefficient between the predicted values and the observed values: $R = r_{\hat{\mathbf{y}}, \mathbf{y}}$. This equation offers another view on the meaning

of the multiple correlation coefficient quantifying the interrelationship (correlation) of the *combined* set of optimally weighted predictor variables with the observed time course.

3.3.3.1 GLM Diagnostics

Note that if the design matrix (model) does not contain all relevant predictors, condition-related increases or decreases in the voxel time course will be explained by the error values instead of the model. It is, therefore, important that the design matrix is constructed with all expected effects, which may also include reference functions not related to experimental conditions, for example, estimated motion parameters or drift predictors if not removed during preprocessing (see Sect. 3.2.2). In case that all expected effects are properly modeled, the residuals should reflect only “pure” noise fluctuations. If some effects are not (correctly) modeled, a plot of the residuals may show low-frequency fluctuations instead of a stationary noise time course. A visualization of the residuals (for selected voxels or regions-of-interest) is, thus, a good diagnostic to assess whether the design matrix has been defined properly.

3.3.3.2 GLM Significance Tests

The multiple correlation coefficient is an important measure of the “goodness of fit” of a GLM. In order to test whether a specified model significantly explains variance in a voxel time course, a F statistic can be calculated for an R value with $p-1$ degrees of freedom in the numerator and $n-p$ degrees of freedom in the denominator:

$$F_{n-1, n-p} = \frac{R^2 (n-p)}{(1-R^2)(p-1)}$$

An error probability value p can then be obtained for the calculated F statistics. A high F value (p value <0.05) indicates that the experimental conditions as a whole have a significant modulatory effect on the data time course (omnibus effect).

While the overall F statistic answers the question whether the specified model significantly explains a voxel’s time course, it does not allow to assess which individual conditions differ significantly from each other. Comparisons between conditions can be formulated as *contrasts*, which are linear combinations of beta values corresponding to null hypotheses. To test, for example, whether activation in a single condition 1 deviates significantly from baseline, the null hypothesis would be that there is no effect in the population, that is, $H_0: b_1=0$. To test whether activation in condition 1 is significantly different from activation in condition 2, the null hypothesis would state that the beta values of the two conditions would not differ: $H_0: b_1=b_2$ or $H_0: (+1)b_1+(-1)b_2=0$. To test whether the mean of condition 1 and condition 2 differs from condition 3, the following contrast could be specified: $H_0: (b_1+b_2)/2=b_3$ or $H_0:$

$(+1)b_1 + (+1)b_2 + (-2)b_3 = 0$. The values used to multiply the respective beta values are often written as a *contrast vector* \mathbf{c} . In the latter example,³ the contrast vector would be written as $\mathbf{c} = [+1 \ +1 \ -2]$. Using matrix notation, the linear combination defining a contrast can be written as the scalar product of contrast vector \mathbf{c} and beta vector \mathbf{b} . The null hypothesis can then be simply described as $\mathbf{c}'\mathbf{b} = 0$. For any number of predictors k , such a contrast can be tested with the following t statistic with $n-p$ degrees of freedom:

$$t = \frac{\mathbf{c}'\mathbf{b}}{\sqrt{\text{Var}(\mathbf{e})\mathbf{c}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{c}}}$$

The numerator of this equation contains the described scalar product of the contrast and beta vector. The denominator defines the standard error of $\mathbf{c}'\mathbf{b}$, that is, the variability of the estimate due to noise fluctuations. The standard error depends on the variance of the residuals $\text{Var}(\mathbf{e})$ as well as on the design matrix \mathbf{X} . With the known degrees of freedom, a t value for a specific contrast can be converted in an error probability value p using the equation shown earlier. Note that the null hypotheses above were formulated as $\mathbf{c}'\mathbf{b} = 0$ implying a *two-sided* alternative hypothesis, that is, $H_a: \mathbf{c}'\mathbf{b} \neq 0$. For one-sided alternative hypotheses, for example, $H_a: b_1 > b_2$, the obtained p value from a two-sided test can be simply divided by 2 to get the p value for the one-sided test. If this p value is smaller than 0.05 and if the t value is positive (since b_1 is assumed to be larger than b_2), the null hypothesis may be rejected.

3.3.3.3 Conjunction Analysis

Experimental research questions often lead to specific hypotheses, which can best be tested by the *conjunction* of two or more contrasts. As an example, it might be interesting to test with contrast \mathbf{c}_1 whether condition 2 leads to significantly higher activity than condition 1 *and* with contrast \mathbf{c}_2 whether condition 3 leads to significantly higher activity than condition 2. This question could be tested with the following conjunction contrast:

$$\mathbf{c}_1 \wedge \mathbf{c}_2 = [-1 \ +1 \ 0] \wedge [0 \ -1 \ +1].$$

Note that a logical “AND” operation is defined for Boolean values (true/false) but that t values associated with individual contrasts can assume any real value. An appropriate way to implement a logical “AND” operation for conjunctions of contrasts with continuous statistical values is to use a mini-

mum operation, that is, the significance level of the conjunction contrast is identical to the significance level of the contrast with the smallest t value: $t_{c_1 \wedge c_2} = \min(t_{c_1}, t_{c_2})$. For more details about conjunction testing, see Nichols et al. (2005).

3.3.3.4 Multicollinear Design Matrices

Multicollinearity exists when predictors of the design matrix are highly intercorrelated. To assess multicollinearity, pair-wise correlations between predictors are not sufficient. A better way to detect multicollinearity is to regress each predictor variable on all the other predictor variables and examine the resulting R^2 values. Perfect or total multicollinearity occurs when a predictor of the design matrix is a linear function of one or more other predictors, that is, when predictors are linearly dependent on each other. While in this case solutions for the GLM system of equations still exist, there is no *unique* solution for the beta values. From a mathematical perspective of the GLM, the square matrix $\mathbf{X}'\mathbf{X}$ becomes *singular*, that is, it loses (at least) one dimension, and is no longer invertible in case that \mathbf{X} exhibits perfect multicollinearity. Matrix inversion is required to calculate the essential term $(\mathbf{X}'\mathbf{X})^{-1}$ used for computing beta values and standard error values (see above). Fortunately, special methods, including singular value decomposition (SVD), allow obtaining (pseudo-) inverses for singular (rank-deficient) matrices. Note, however, that in this case the absolute values of beta weights may be difficult to interpret, and statistical hypothesis tests must meet special restrictions.

In fMRI design matrices, multicollinearity occurs if all conditions are modeled as predictors in the design matrix including the baseline (rest, control) condition. Without the baseline condition, multicollinearity is avoided and beta weights are obtained which are easily interpretable. As an example consider the case of two main conditions and a rest condition. If we would not include the rest condition (recommended), the design matrix would not be multicollinear and the two beta weights b_1 and b_2 would be interpretable as increase or decrease of activity relative to the baseline signal level modeled by the constant term (Fig. 17, right). Contrasts could be specified to test single beta weights, for example, the contrast $\mathbf{c} = [1 \ 0]$ would test whether condition 1 leads to significant (de)activation. Furthermore, the two main conditions could be compared with the contrast $\mathbf{c} = [-1 \ 1]$, which would test whether condition 2 leads to significantly more activation than condition 1. If the design matrix would include a predictor for the rest condition, we would obtain perfect multicollinearity and the matrix $\mathbf{X}'\mathbf{X}$ would be singular. Using a pseudo-inverse or SVD approach, we would obtain now three beta values (plus the constant), one for the rest condition, one for main condition 1, and one for main

³Note that the constant term is treated as a confound and it is not included in contrast vectors, i.e., it is implicitly assumed that b_0 is multiplied by 0 in all contrasts. To include the constant explicitly, each contrast vector must be expanded by one entry at the beginning or end.

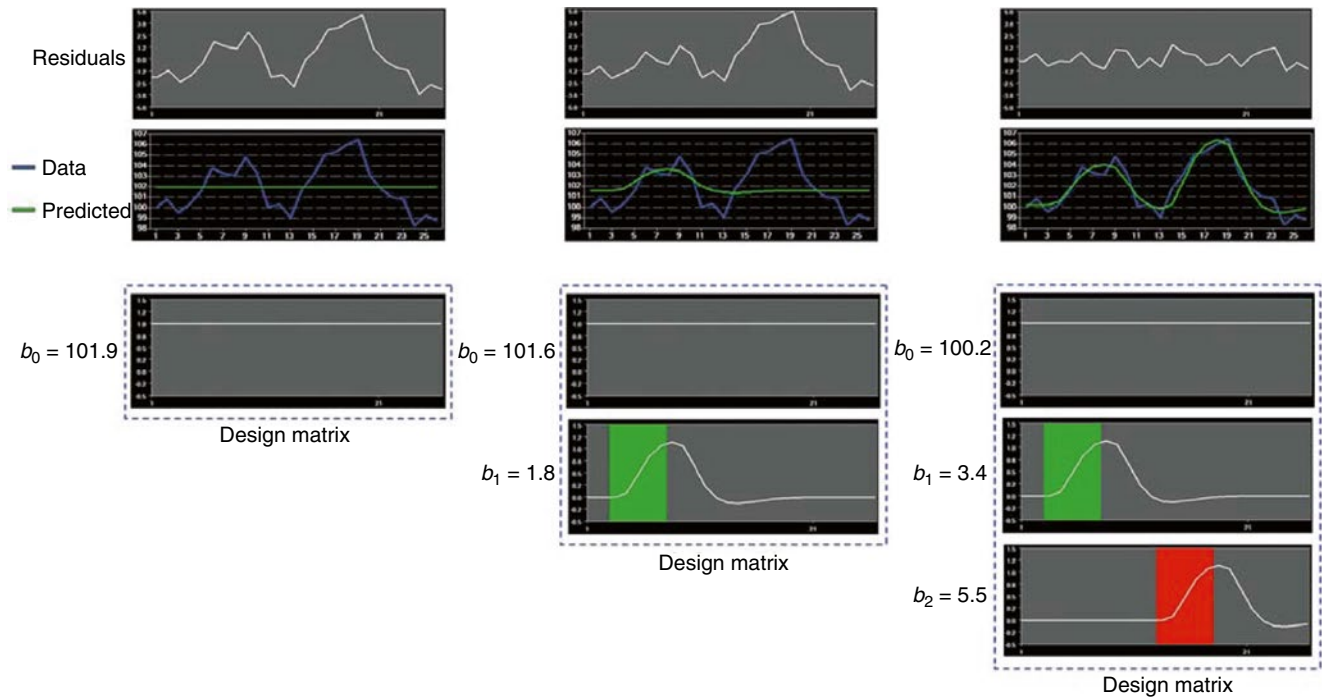


Fig. 17 Three GLMs fitting the same data with different design matrices. *Top row* shows residuals, *second row* predicted (green) and observed (blue) voxel time course. The design matrix on the left contains only one predictor, the constant term. The estimated beta weight (b_0) scales the constant term to the mean signal level. The design matrix in the middle adds a predictor for green main condition. The estimated beta weights (b_0, b_1) scale the predictors and the

weighted sum explains more variance than first model, but residual variance is still high. The third model (*right*) adds a predictor for red main condition. The estimated beta weights (b_0, b_1, b_2) scale the predictors and weighted sum now explains all task-related signal fluctuations. The residuals reflect now only noise. The example highlights the importance of modeling all known effects in the design matrix

condition 2. While the values of beta weights might not be interpretable, correct inferences of contrasts can be obtained if an additional restriction is met, typically that the sum of the contrast coefficients equals 0. To test whether main condition 1 differs significantly from the rest condition, the contrast $\mathbf{c}=[-1+1\ 0]$ would now be used. The contrast $\mathbf{c}=[0-1+1]$ would be used to test whether condition 2 leads to more activation than condition 1.

3.3.3.5 GLM Assumptions

Given a correct model (design matrix), the standard estimation procedure of the GLM – ordinary least squares (OLS) – operates correctly only under the following assumptions. The population error values ε must have an expected value of zero and constant variance at each time point i :

$$E[\varepsilon_i] = 0$$

$$\text{Var}[\varepsilon_i] = \sigma^2$$

Furthermore, the error values are assumed to be uncorrelated:

$$\text{cov}(\varepsilon_i, \varepsilon_j) = 0 \text{ for all } i \neq j$$

To justify the use of t and F distributions in hypothesis tests, errors are further assumed to be normally distributed:

$$\varepsilon_i \sim N(0, \sigma^2)$$

In summary, errors are assumed to be normal independent and identically distributed (often abbreviated as “normal i.i.d.”). Under these assumptions, the solution obtained by the least squares method is optimal in the sense that it provides the most efficient unbiased estimation of the beta values. While the OLS approach is robust with respect to small violations, assumptions should be checked. In the context of fMRI measurements, the assumption of uncorrelated error values requires special attention.

3.3.3.6 Correction for Serial Correlations

In fMRI data, one typically observes *serial correlations*, that is, high values are followed more likely by high values than low values and vice versa. The assessment of these serial correlations is not performed on the original voxel time course but on the time course of the residuals since serial correlations in the recorded signal are expected to some

extent from slow task-related fluctuations. Task-unrelated serial correlations most likely occur because data points are measured in rapid succession, that is, they are also observed when scanning phantoms. Likely sources of temporal correlations are physical and physiological noise components such as hardware-related low-frequency drifts, oscillatory fluctuations related to respiration and cardiac pulsation, and residual head motion artifacts. Serial correlations violate the assumption of uncorrelated errors (see section above). Fortunately the beta values estimated by the GLM are correct estimates (unbiased) even in presence of serial correlations. The standard errors of the betas are biased, however, leading to “inflated” test statistics, that is, t or F values are higher than they should be. This can be explained by considering that the presence of serial correlations (serial dependence) reduces the true number of independent observations (effective degrees of freedom) that will, thus, be lower than the nominal number of observations. Without correction, the degrees of freedom are systematically overestimated leading to an underestimation of the error variance resulting in inflated statistical values, that is, t or F values are too high. It is, thus, necessary to correct for serial correlations in order to obtain valid error probabilities. Serial correlations can be corrected using several approaches. In *pre-whitening* approaches, autocorrelation is first estimated and removed from the data; the pre-whitened data can then be analyzed with a standard OLS GLM solution. In *pre-coloring* approaches (e.g., Friston et al. 1995), a strong autocorrelation structure is imposed on the data by temporal smoothing and degrees of freedom are adjusted according to the imposed (known) autocorrelation. The pre-coloring (temporal smoothing) operation acts, however, as a low-pass filter and may weaken experimentally induced signals of interest and is thus not the preferred method. The pre-whitening approach can be expressed in terms of a more powerful estimation procedure than OLS called *generalized least squares* (GLS, Searle et al. 1992). As opposed to the OLS method, GLS works correctly also in case that error values exhibit correlations or when error variances are not homogeneous. Note, however, that this more powerful estimation approach only provides correct results in case that the true (population) variances and covariances of the error values are known. With the known error covariance matrix \mathbf{V} , the betas and their (co-)variances can be calculated with GLS as follows:

$$\mathbf{b} = (\mathbf{X}'\mathbf{V}^{-1}\mathbf{X})^{-1} \mathbf{X}'\mathbf{V}^{-1} \mathbf{y}$$

$$\text{cov}(\mathbf{b}) = (\mathbf{X}'\mathbf{V}^{-1}\mathbf{X})^{-1}$$

With the obtained b values and their covariances, any contrast can then be assessed statistically as described above for the OLS method. When comparing the GLS solution with the OLS solution, it is evident that the inverse of the population error

covariance matrix \mathbf{V}^{-1} is needed to properly treat the effect of covariance of the errors on the parameter estimates (betas and their covariances). Note also that when setting \mathbf{V} as a diagonal matrix (entries outside the main diagonal are zero, i.e., no covariation of errors) with equal variance values (all values of the main diagonal are the same, e.g., 1), the GLS equation reduces to the OLS solution, that is, the \mathbf{V}^{-1} term vanishes.

Since the population covariance matrix of the error values \mathbf{V} is usually not known, it needs to be estimated from the data itself. Since there are too many degrees of freedom (number of time points squared: n^2), \mathbf{V} cannot be estimated for the general case of arbitrary covariance matrices. It is, however, often possible to estimate \mathbf{V} for special cases where only some parameters need to be estimated. The two most important special cases in the context of fMRI data analysis are the treatment of serial correlations (see below) and the treatment of unequal variances when integrating data from different subjects in the context of mixed-effects group analyses.

A simple pre-whitening procedure was developed (Cochrane and Orcutt 1949; Bullmore et al. 1996) independently from the GLS approach but can be shown to be identical to a GLS solution. The procedure assumes that the errors follow a first-order autoregressive, or $AR(1)$, process. After calculation of a GLM using OLS, the amount of serial correlation a_1 is estimated using pairs of successive residual values (e_t, e_{t+1}), that is, the residual time course is correlated with itself shifted by one time point (lag = 1). In the second step, the estimated serial correlation is removed from the measured voxel time course by calculating the transformed time course $y_t^n = y_{t+1} - a_1 \cdot y_t$. The superscript “ n ” indicates the values of the new, adjusted time course. The same calculation is also applied to each predictor time course resulting in an adjusted design matrix \mathbf{X}^n . In the third step, the GLM is recomputed using the adjusted voxel time course and adjusted design matrix resulting in correct standard errors for beta estimates and, thus, correct significance levels for contrasts (of course under the assumption that the $AR(1)$ model is correct). If autocorrelation is not sufficiently reduced in the new residuals, the procedure can be repeated. If performed using the GLS approach, the first step is identical to the Cochrane-Orcutt method, that is, OLS is used to fit the GLM and the obtained residuals are used to estimate the value of the $AR(1)$ term. The adjustment of the time course y_t and the design matrix described above need, however, not be performed explicitly since these adjustments are handled implicitly in the next step by using a \mathbf{V}^{-1} term in the GLS equations that contains values in the off-diagonal elements derived from the estimated serial correlation term.

While an $AR(1)$ autocorrelation model substantially reduces serial correlations in fMRI data, better results are obtained when using an $AR(2)$ model, that is, both first-order and second-order autocorrelation terms should be estimated and used to construct the error covariance matrix \mathbf{V} for GLS estimation. Since serial correlations differ across voxels,




Fig. 18 Comparison of two methods used to solve the multiple comparisons problem. A statistical map has been computed comparing responses to faces and houses. *Red/yellow colors* depict regions with larger responses to faces than to houses, while *blue regions* indicate areas with larger responses to houses than to faces. (a) No correction for multiple comparisons has been performed. (b)

Thresholding result when using the false discovery rate approach (FDR). (c) Thresholding result when using the Bonferroni method. The p values shown on *top* of each panel have been used to threshold the map as provided by the respective method. The FDR method shows more voxels as significant because it is less conservative than the Bonferroni method

serial correlation correction should be performed separately for each voxel time course as opposed to the (also used) estimation of serial correlation values from multiple averaged (neighboring) voxel time courses. An AR(2) serial correlation model applied separately for each voxel time course has been recently shown to be the most accurate approach to treat serial correlations when compared to other models (Lenoski et al. 2008).

3.3.4 Creation of Statistical Maps

The statistical analysis steps were described for a single voxel's time course since standard statistical methods are performed independently for each voxel (univariate "voxel-wise" analysis). Since a typical fMRI data set contains several hundred thousand voxels, a statistical analysis is performed independently hundred thousands of times. Running a GLM, for example, results in a set of estimated beta values attached to each voxel. A specified contrast $\mathbf{c}'\mathbf{b}$, will be performed using the same contrast vector \mathbf{c} for each voxel v , but it will use a voxel's vector of beta values \mathbf{b}_v (and the voxel's error term) to obtain voxel-specific t and p values. Statistical test results for individual voxels are integrated in a 3D data set called a *statistical map*. To visualize a statistical map, the obtained values, for example, contrast t values, can be shown at the location of each voxel replacing anatomical intensity values shown as default. As a further useful condition, the statistical values are often only shown for those voxels exceeding a specified *statistical threshold*. This allows visualizing anatomical information in large parts of the brain while statistical information is shown (overlaid) only in those regions exhibiting suprathreshold (usually statistically significant) signal modulations. While anatomical information is normally visualized using a range of gray values, suprathreshold statistical test values are typically visualized using multiple colors, for example, a red-to-yellow range for positive values and a green-to-blue range for negative values. With these colors, a positive (negative) t value just passing a

specified threshold would be colored in red (green), while a very high positive (negative) t value would be colored in yellow (blue) (Fig. 18).

3.3.5 The Multiple Comparison Problem

An important issue in fMRI data analysis is the specification of an appropriate threshold for statistical maps. If there would be only a single voxel's data, a conventional threshold of $p < 0.05$ (or $p < 0.01$) could be used to assess significance of an observed effect quantified by an R , t , or F statistic. Running the statistical analysis separately for each voxel creates, however, a *massive multiple comparison problem*. If a *single test* is performed, the conventional threshold protects from wrongly declaring a voxel as significantly modulated (false positive) with a probability of $p < 0.05$ when there is no effect in the population (α error). Note that in case that the null hypothesis (no effect) holds, an adopted error probability of $p = 0.05$ implies that if the same test would be repeated 100 times, the alternative hypothesis would be accepted wrongly on average in five cases, that is, we would expect 5 % of false positives. If we assume that there is no real effect in any voxel time course, running a statistical test spatially in parallel is statistically identical to repeating the test 100,000 times at a single voxel (each time with new measured data). It is evident that this would lead to about 5,000 false positives, that is, about 5,000 voxels would be labeled "significant" although these voxels would reach the 0.05 threshold purely due to chance.

Several methods have been suggested to control this massive multiple comparison problem. The *Bonferroni correction* method is a simple multiple comparison correction that controls the α error *across all voxels*, and it is therefore called a *family-wise error* (FWE) correction approach. The method calculates single-voxel threshold values in such a way that an error probability of 0.05 is obtained at the global level. With N independent tests, this is achieved by using a statistical

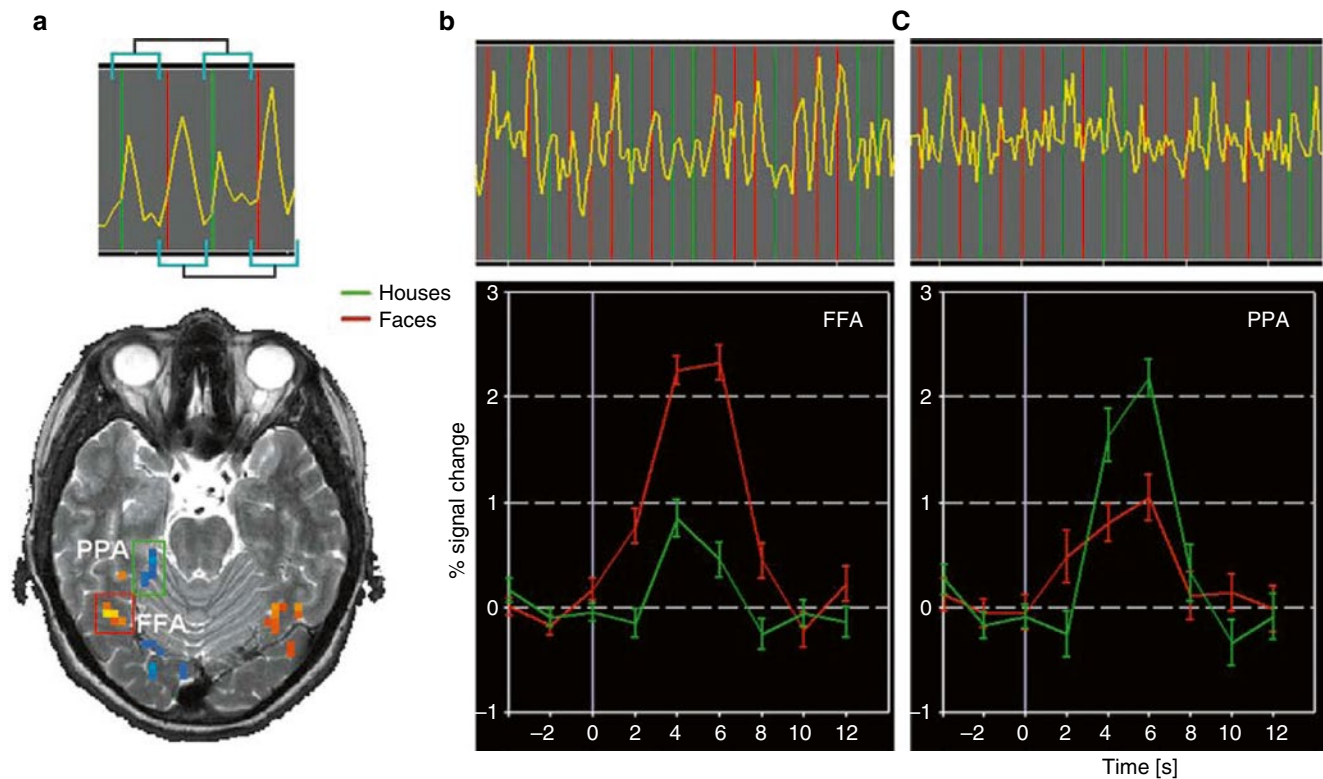


Fig. 19 Principle of event-related averaging and event-related averaging plots from a slow event-related design. (a) The thresholded statistical map shows in *red/yellow color* brain regions responding more to faces than to houses and in *blue color* brain regions responding more to houses than to faces. The areas demarcated with *red and green rectangles* in the lower panel correspond well to *fusiform face*

area (FFA) and parahippocampal place area (PPA), respectively (O’Craven and Kanwisher 2000). (b) Time course from FFA (upper panel) and event-related averaging plot (lower panel) obtained by selectively averaging all responses belonging to the same condition. (c) Time course (upper panel) and event-related averaging plot (lower panel) from PPA

significance level which is N times smaller than usual. The Bonferroni correction can be derived mathematically as follows. Under the assumption of independent tests, the probability that all of N performed tests lead to a sub-threshold result is $(1-p)^N$ and the probability to obtain one or more false-positive results is $1-(1-p)^N$. In order to guarantee a family-wise (global) error probability of $p_{FWE} = 1-(1-p)^N$, the threshold for a single test, p , has to be adjusted as follows: $p = 1 - (1 - p_{FWE})^{1/N}$. For small p_{FWE} values (e.g., 0.05), this equation can be approximated by $p = p_{FWE}/N$. This means that to obtain a global error probability of $p_{FWE} < 0.05$, the significance level for a single test is obtained by dividing the family-wise error probability by the number of independent tests. Given 100,000 voxels, we would obtain an adjusted single-voxel threshold of $p_v = p_{FWE}/N = 0.05/100,000 = 0.0000005$. The Bonferroni correction method ensures that we do not declare even a single voxel wrongly as significantly activated with an error probability of 0.05. For fMRI data, the Bonferroni method would be a valid approach to correct the α error if the data at neighboring voxels would be truly independent from each other. Neighboring voxels, however, show similar response patterns within functionally

defined brain regions, such as the fusiform face area (FFA). In the presence of such spatial correlations, the Bonferroni correction method operates too conservative, that is, it corrects the error probability more strongly than necessary. As a result of a too strict control of the α error, the sensitivity (power) to detect truly active voxels is reduced: Many voxels will be labeled as “not significant” although they likely reflect true effects. As described earlier, wrongly accepting (rejecting) a null (alternative) hypothesis is called type II error or β error.

Worsley et al. (1992) suggested a less conservative approach to correct for multiple comparisons taking explicitly the observation into account that neighboring voxels are not activated independently from each other but are more likely to activate together in clusters. In order to incorporate spatial neighborhood relationships in the calculation of global error probabilities, the method describes a statistical map as a *Gaussian random field* (for details, see Worsley et al. 1992). Unfortunately, application of this correction method requires that the fMRI data are spatially smoothed substantially reducing one of its most attractive properties, namely, the achievable high spatial resolution.

Another correction method incorporating the observation that neighboring voxels often activate in clusters is based on Monte Carlo simulations that generate many random images (maps) using the spatial correlation structure of the original map; the generated maps are used to calculate the likelihood to obtain different sizes of functional clusters by chance for specific (less conservative) single-voxel thresholds (Forman et al. 1995). The calculated cluster extent threshold combined with a less strict single-voxel threshold is finally applied to the statistical map ensuring that a global error probability of $p < 0.05$ is met. This approach does not require spatial smoothing and appears highly appropriate for fMRI data. A disadvantage is that the method is quite compute intensive and that small functional clusters might not be discovered.

While the described multiple comparison correction methods aim to control the family-wise error rate, the *false discovery rate (FDR) approach* (Benjamini and Hochberg 1995) uses a different statistical logic and has been proposed for fMRI analysis by Genovese and colleagues (2002). This approach does not control the overall number of false-positive voxels but the number of false-positive voxels among the subset of voxels labeled as significant. Given a specific threshold, suprathreshold voxels are called “discovered” voxels or “voxels declared as active.” With a specified false discovery rate of $q < 0.05$, one would accept that 5 % of the discovered (suprathreshold) voxels would be false positives. Given a desired false discovery rate, the FDR algorithm calculates a single-voxel threshold, which ensures that the voxels beyond that threshold contain on average not more than the specified proportion of false positives. With a q value of 0.05, this also means that one can “trust” 95 % of the suprathreshold (i.e., color-coded) voxels since the null hypothesis has been rejected correctly. Since the FDR logic relates the number of false positives to the amount of truly active voxels, the FDR method adapts to the amount of activity in the data: The method is very strict if there is not much evoked activity in the data but assumes less conservative thresholds if a larger number of voxels show task-related effects. In the extreme case that not a single voxel is truly active, the calculated single-voxel threshold is identical to the one computed with the Bonferroni method. The FDR method appears ideal for fMRI data because it does not require spatial smoothing and it detects voxels with a high sensitivity (low β error) if there are true effects in the data.

Another simple approach to the multiple comparisons problem is to reduce the number of tests by using anatomical masking. Most correction methods, including Bonferroni and FDR, can be combined with this approach since a smaller number of tests leads to a less strict control of the α error and thus a smaller β error is made as compared to inclusion of all

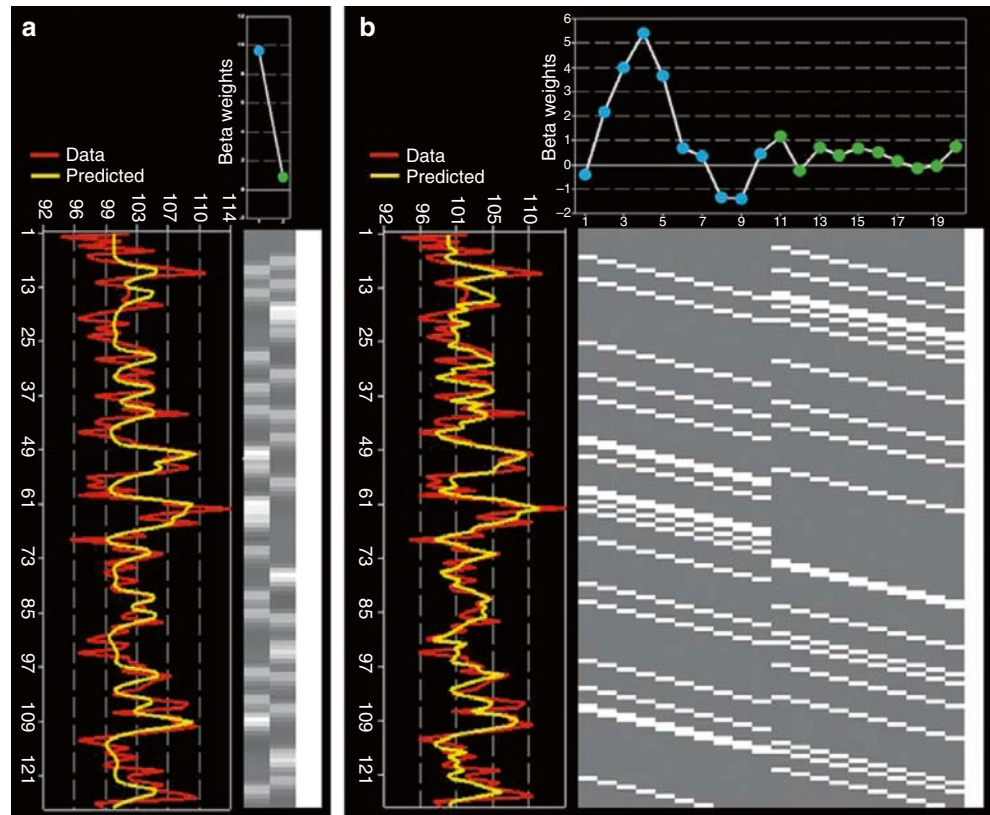
voxels. In a simple version of an anatomical mask, an intensity threshold for the basic signal level can be used to remove voxels outside the head. The number of voxels can be further reduced by masking the brain, for example, after performing a brain extraction step. These simple steps typically reduce the number of voxels from about 100,000 to about 50,000 voxels. In a more advanced version (Goebel and Singer 1999) statistical data analysis may be restricted to gray matter voxels, which may be identified by standard cortex segmentation procedures (e.g., Kriegeskorte and Goebel 2001). This approach not only removes voxels outside the brain but also excludes voxels in the white matter and ventricles. Note that anatomically informed correction methods do not require spatial smoothing of the data and not only reduce the multiple comparisons problem but also reduce computation time since fewer tests (e.g., GLM calculations) have to be performed.

3.3.6 Event-Related Averaging

Event-related designs cannot only be used to *detect* activation effects but also to *estimate* the time course of task-related responses. Visualization of mean response profiles can be achieved by averaging all responses of the same condition across corresponding time points with respect to stimulus onset. Averaged (or even single trial) responses can be used to characterize the temporal dynamics of brain activity within and across brain areas by comparing estimated features such as response latency, duration, and amplitude (e.g., Kruggel and von Cramon 1999; Formisano and Goebel 2003). In more complex, temporally extended tasks, responses to subprocesses may be identified. In working memory paradigms, for example, encoding, delay, and response phases of a trial may be separated. Note that event-related selective averaging works well only for slow event-related designs. In rapid event-related designs, responses from different conditions lead to substantial overlap, and event-related averages are often meaningless. In this case, deconvolution analysis is recommended (see below).

In order to avoid circularity, event-related averages should only be used descriptively if they are selected from significant clusters identified from a whole-brain statistical analysis of the same data. Even a merely descriptive analysis visualizing averaged condition responses is, however, helpful in order to ensure that significant effects are caused by “BOLD-like” response shapes and not by, for example, signal drifts or measurement artifacts. If ROIs are determined using independent (localizer) data, event-related averages extracted from these regions in a subsequent (main) experiment can be statistically analyzed. For a more general discussion of ROI vs. whole-brain analyses, see Friston and Henson (2006), Friston et al. (2006), Saxe et al. (2006), and Frost and Goebel (2013).

Fig. 20 Deconvolution analysis of a rapid event-related design. Time runs from *top to bottom*, design matrix depicted in graphical view. Beta values are plotted horizontally at positions corresponding to the respective predictor. **(a)** Standard analysis with two main predictors obtained by convolution of stimulus times with standard hemodynamic response model (two gamma function). Beta values can be compared with a standard contrast. **(b)** Deconvolution analysis fitting the same data. Each condition is modeled with ten “stick” predictors allowing to estimate the time course of condition-related responses as if stimuli were presented in a slow event-related design. Beta values may be compared within and across conditions



3.3.7 Deconvolution Analysis

While standard design matrix construction (convolution of box-car with two gamma function) can be used to estimate condition amplitudes (beta values) in rapid event-related designs, results critically depend on the appropriateness of the assumed standard BOLD response shape: Due to variability in different brain areas within and across subjects, a *static* model of the response shape might lead to non-optimal fits. Furthermore, the isolated responses to different conditions cannot be visualized due to overlap of condition responses over time. To model the shape of the hemodynamic response more flexibly, multiple basis functions (predictors) may be defined for each condition instead of a single predictor. Two often-used additional basis functions are derivatives of the two gamma function with respect to two of its parameters, delay and dispersion. If added to the design matrix for each condition, these basis functions allow capturing of small variations in response latency and width of the response. Other sets of basis functions (i.e., gamma basis set, Fourier basis set) are much more flexible, but obtained results are often more difficult to interpret. *Deconvolution analysis* is a general approach to estimate condition-related response profiles using a flexible and interpretable set of basis functions. It can be easily implemented as a GLM by defining an appropriate design matrix (Fig. 20) that models each bin after

stimulus onset by a separate condition predictor (delta or “stick” functions). This is also called a finite impulse response (FIR) model because it allows estimating any response shape evoked by a short stimulus (impulse). In order to capture the BOLD response for short events, about 20 s is typically modeled after stimulus onset. This would require, for each condition, 20 predictors in case of a TR of 1 s or ten predictors in case of a TR of 2 s (Fig. 20). Despite overlapping responses, fitting such a GLM “recovers” the underlying condition-specific response profiles in a series of beta values, which appear in plots as if event-related averages have been computed in a slow event-related design (Fig. 20). Since each condition is modeled by a series of temporally shifted predictors, hypothesis tests can be performed that compare response amplitudes *at different moments in time* within and between conditions. Note, however, that the deconvolution analysis assumes a linear time invariant system (see Sect. 3.1). In order to uniquely estimate the large number of beta values from overlapping responses, variable ITIs must be used in the experimental design (see Sect. 3.1). The deconvolution model is very flexible allowing to capture any response shape. This implies that also non-BOLD-like time courses will be detected easily since the trial responses are not “filtered” by the ideal BOLD response shape as in conventional analysis.

3.4 Integration of Anatomical and Functional Data

The localization of the neural correlates of sensory, motor, and cognitive functions requires a precise relationship between voxels in calculated statistical maps with voxels in high-resolution anatomical data sets. This is especially important in single-subject analyses and, thus, for presurgical mapping. While it is recommended to also view statistical maps overlaid on a volume of the functional data itself, EPI data sets often do not contain sufficient anatomical detail to specify the precise location of an active cluster in a subject's brain. 3D renderings of high-resolution anatomical data sets may greatly aid in visualizing activated brain regions. Advanced visualization requires that a high-resolution 3D data set is recorded for a subject and that the functional data is *coregistered* to the 3D data set as precisely as possible. Anatomical data sets are also important for most brain normalization methods, which is a prerequisite of the analysis of whole-brain group studies. High-resolution anatomical data sets are typically recorded with T_1 -weighted MRI sequences. A typical structural scan covering the whole brain with a resolution of 1 mm in all three dimensions (e.g., 180 sagittal slices) lasts between 5 and 20 min on current 1.5 and 3.0 scanners.

3.4.1 Visualizing Statistical Maps on Anatomical Images

Having identified a statistically significant region in the functional data set does not easily allow a precise statement about its location in the brain of the subject since the functional data itself does often not contain enough anatomical details.⁴ If anatomical, coplanar images are available, it is already helpful to overlay the functional results (thresholded statistical maps) on these “in-plane” images. Figure 19a shows, for example, a statistical map on a high-resolution, coplanar, T_2 -weighted image. While high-resolution, coplanar images improve localization within the recording plane, the direction along slices is sampled with low resolution due to typical distances between slices of 3–5 mm (slice distance = slice thickness + slice gap). Identification of the anatomical substrate of an activated cluster greatly benefits from visualizing functional data over isotropic high-resolution 3D data sets. Overlaying or fusing images from functional data (MRI, PET, SPECT) with high-resolution anatomical MRI data sets is a common visualization method in functional

imaging. In order to correctly fuse functional and anatomical data sets, appropriate coregistration transformations have to be performed.

3.4.2 Coregistration of Functional and Anatomical Data Sets

If functional images are superimposed on coplanar images, spatial transformations (translations and rotations) to align the two data sets are not necessary (except maybe the correction of small head movements and small geometric distortions), since the respective slices are measured at the same 3D positions. Since the coplanar anatomical images are usually recorded with a higher resolution (typically with a 256×256 matrix) than the functional images (typically 64×64 or 128×128 matrices), only a scaling factor has to be applied. To allow high-quality visualization of the functional data in arbitrary resliced anatomical planes, the functional data must be coregistered with high-resolution 3D data sets.

These high-resolution 3D data sets are usually recorded with different slice orientation and position than the functional data, and the coregistration step, thus, requires an affine spatial transformation including translation, rotation, and scaling. These three elementary spatial transformations can be integrated in a single transformation step expressed in a standard 4×4 spatial transformation matrix. If the high-resolution 3D data set has been recorded in the same scanning session as the functional data, the coregistration matrix can be constructed simply by using the scanning parameters (slice positions, pixel resolution, slice thickness) from both recordings. The alignment based on this information would be perfect if there would be no head movement between the anatomical and functional images. To further improve coregistration results, an additional intensity- or gradient-driven alignment step is usually performed after the initial (mathematical) alignment correcting for head displacements (and eventually geometric distortions) between the functional and anatomical recordings. While this step operates similar as described for motion correction, it is likely not possible to align the two data sets perfectly well in all regions of the brain due to signal dropouts and distortions in the functional EPI images. For neurosurgical purposes, it is important to ensure that at least the relevant regions of the brain do not suffer from EPI distortions and that they are perfectly coregistered with the anatomical data. EPI distortions and signal dropouts can be corrected to some extent with special MRI sequence modifications as well as with image processing software. Using appropriate visualization tools, it is also possible to manually align a functional volume with an anatomical 3D data set. The precision of manual alignment depends, however, strongly on acquired expertise.

⁴While functional sequences are T_2^* weighted, the first functional volume of a run contains the richest anatomical detail because it is T_1 weighted. Unfortunately, this data set is often thrown away either by the scanner directly (during “prep scans”) or by transfer of the data to the researcher. We recommend to keep the first functional volume and to use it for visualization and coregistration because of its relative richness in anatomical details.




Fig. 21 Cortex representations used for advanced visualization. (a) Segmentation and surface reconstruction of the inner (*white/gray matter, yellow*) and outer (*gray matter/CSF, magenta*) boundary of the

gray matter. (b) “Inflated” cortex representation of the left hemisphere obtained by iterative morphing process. (c) “Flat map” of the right cortical hemisphere with superimposed functional data

3.4.3 Visualizing Statistical Maps on Reconstructed Cortex Representations

High-resolution anatomical data sets can be used to create 3D volume or surface renderings of the brain, which allow additional helpful visualizations of functional data on a subject’s brain (Fig. 21c). These visualizations require segmentation of the brain, which can be performed automatically with most available software packages. For more advanced visualizations, segmentation of cortical voxels allows to construct topologically correct mesh representations of the cortical sheet, one for the left and one of the right hemisphere (e.g., Fischl et al. 1999; Kriegeskorte and Goebel 2001). The obtained meshes (Fig. 21a) may be further transformed into inflated (Fig. 21b) and flattened (Fig. 21c) cortex representations. Functional data can then be superimposed on folded, inflated, and flattened representations (Fig. 21c), which is particularly useful for topologically organized functional information, for example, in the context of retinotopic, tonotopic, and somatotopic mapping experiments. To help in orientation, inflated and flattened cortex representations indicate gyral and sulcal regions by color-coding local curvature; concave regions, indicating sulci, may be depicted, for example, with a dark gray color, while convex regions, indicating gyri, may be depicted, for example, with a light gray color (Fig. 21). A general advantage of visualizing functional data on flat maps is that all cortical activation foci from different experiments can be visualized at once at their correct anatomical location in a canonical view. In contrast, visualizing several activated regions using a multi-slice representation

depends on the chosen slice orientation and number of slices. Note that anatomical data is not only important to visualize functional data. Anatomical information may also be used to constrain statistical data analysis as has been described in Sect. 3.3.5. Furthermore, the explicit segmentation of cortical voxels is also the prerequisite for advanced anatomical analyses, including cortical thickness analysis.

3.5 Group Analysis of Functional Data Sets

Presurgical neuroimaging requires detailed single-subject analyses, which can be performed with the methods described in the previous sections. A standardized routine for analyzing (clinical) fMRI data in individuals is given in chapter [Task-based presurgical functional MRI in patients with brain tumors](#). If, however, characterization and statistical assessment of general brain patterns is desired, multiple subjects have to be integrated in groups. Such *group studies* allow generalizing findings from a sample of subjects to the population from which the patients or healthy subjects have been drawn. Group analysis of functional data sets is of clinical relevance when the effects of various brain pathologies or different therapies (e.g., pharmacological effects) on brain function are subject to study.

The integration of fMRI data from multiple subjects is challenging because of the *spatial correspondence problem* between different brains. This problem manifests itself already at a purely anatomical level but presents a funda-

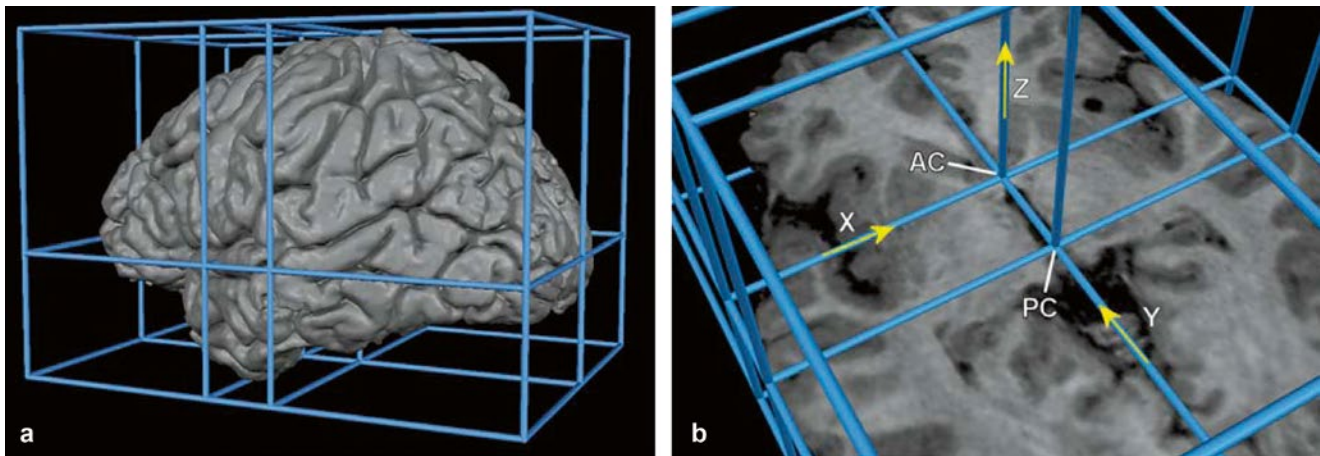


Fig. 22 Definition of Talairach space. (a) View from *left*. (b) View from *top*. Talairach space is defined by three orthogonal axes pointing from left to right (x -axis), posterior to anterior (y -axis), and inferior to superior (z -axis). The origin of the coordinate system is defined by the anterior commissure (AC). Coordinates are in millimeters. The posterior

commissure (PC) is located on the y -axis ($y=-23$ mm). The borders of the Talairach grid (a) correspond to the borders of the cerebrum. The most right point of the brain corresponds to $x=68$ mm, the most left one to $x=-68$ mm, the most anterior one to $y=70$, the most posterior one to $y=-102$, the most upper one to $z=74$, and the most lower one to $z=-42$

mental problem of neuroscience when considered as a question of the consistency of structure-function relationships. At the anatomical level, the correspondence problem refers to the differences in brain shape and, more specifically, to differences in the gyral and sulcal pattern varying substantially across subjects. At this macroanatomical level, the correspondence problem would be solved if brains could be matched in such a way that for each macroanatomical structure in one brain, the corresponding region in the other brain would be known. In neuroimaging, the matching of brains is usually performed by a process called *brain normalization*, which involves warping each brain into a common space allowing averaging over (more or less) corresponding regions in different subjects. After brain normalization, a point in the common space identified by its x , y , and z coordinates is assumed to refer to a similar region in any other normalized brain. The most commonly used target space for normalization is the *Talairach space* (see below) and the closely related *MNI template space*. Unfortunately warping brains in a common space does not solve the anatomical correspondence problem very well, that is, macroanatomical structures, such as banks of prominent sulci are often still misaligned with deviations in the order of 0.5–1 cm. In order to increase the chance that corresponding regions overlap, functional data is therefore often smoothed with a Gaussian kernel with a width of about 1 cm. More advanced anatomical matching schemes attempt to directly align macroanatomical structures such as gyri and sulci (see below) and require less (or no) spatial smoothing of functional data.

The deeper version of the correspondence problem addresses the fundamental question of the existence of an identical relationship between certain brain functions and

neuroanatomical structures across subjects. While neuroimaging has successfully demonstrated that there is common structure-function relationship across brains, a high level of variability has also been observed, especially for higher cognitive functions. A more satisfying answer to this fundamental question might only emerge after much more careful investigations, for example, by letting the same subjects perform a large battery of tasks (Frost and Goebel 2012). An interesting approach to the functional correspondence problem has been proposed that aims to only align those brain regions-of-interest (ROIs), which are activated in a given task in all or most subjects.

3.5.1 Talairach Transformation

The most often-used standard space for brain normalization is the *Talairach space* (Talairach and Tournoux 1988) or the closely related *MNI template space*. Talairach transformation is controlled either by the (automatic) specification of a few prominent landmarks or by a data-driven alignment of a subject's brain to a target (average) brain (typically the MNI template brain) that has been previously transformed in (near-) Talairach space. In the explicit landmark-based approach (Talairach and Tournoux 1988), the midpoint of the anterior commissure (AC) is located first, serving as the origin of Talairach space. The brain is then rotated around the new origin (AC) so that the posterior commissure (PC) appears in the same axial plane as the anterior commissure (Fig. 22). The connection of AC and PC in the middle of the brain forms the y -axis of the Talairach coordinate system. The x -axis runs from the left to the right hemisphere through AC orthogonal to the y -axis. The z -axis runs from the inferior part of the brain to the superior part through AC orthogonal to both other axes. In order to further constrain the x - and

z -axes, a y - z plane is rotated around the y (AC-PC)-axis until it separates the left and right hemisphere (midsagittal plane). The obtained AC-PC space is attractive for individual clinical applications, especially presurgical mapping and neuro-navigation since it keeps the original size of the subject's brain intact while providing a common orientation for each brain anchored at important landmarks. For a full Talairach transformation, a cuboid is defined running parallel to the three axes enclosing precisely the cortex. This cuboid or bounding box requires specification of additional landmarks specifying the borders of the cerebrum. The bounding box is subdivided by several sub-planes. The midsagittal y - z plane separates two sub-cuboids containing the left and right hemisphere, respectively. An axial (x - y) plane through the origin separates two sub-cuboids containing the space below and above the AC-PC plane. Two coronal (x - z) planes, one running through AC and one running through PC, separate three sub-cuboids: the first contains the anterior portion of the brain anterior to the AC, the second contains the space between AC and PC, and the third contains the space posterior to PC. These planes separate 12 sub-cuboids. In a final Talairach transformation step, each of the 12 sub-cuboids is expanded or shrunken linearly to match the size of the corresponding sub-cuboid of the standard Talairach brain. To reference any point in the brain, x , y , and z coordinates are specified in millimeters of Talairach space. Talairach and Tournoux (1988) also defined the "proportional grid" to reference points within the defined cuboids.

In summary, Talairach normalization ensures that the anterior and posterior commissures obtain the same coordinates in each brain and that the sub-cuboids defined by the AC-PC points and the borders of the cortex will have the same size. Note that the specific distances between landmarks in the original postmortem brain are not important for establishing the described spatial relationship between brains. The important aspect of Talairach transformation is that correspondence is established across brains by linearly interpolating the space between important landmarks.

While Talairach transformation provides a recipe to normalize brains, regions at the same coordinates in different individuals do not necessarily point to homologous brain areas. This holds especially true for cortical regions (e.g., Frost and Goebel 2012). For subcortical structures around the AC-PC landmarks, however, the established correspondence is remarkably good even when analyzing high-resolution fMRI data (e.g., De Martino et al. 2013).

As an alternative to specify crucial landmarks, a direct approach of stereotactic normalization has been proposed (e.g., Evans et al. 1993; Ashburner and Friston 1999) that attempts to align each individual brain as good as possible to an average target brain, called *template brain*. The most often-used template brain is provided by the Montréal Neurological Institute (MNI) and has been created by aver-

aging many (>100) single brains after manual Talairach transformation. Although automatic alignment to a template brain has the potential to result in a better correspondence between brain regions, comparisons have shown that the achieved results are not substantially improved as compared to the explicit landmark specification approach, even when using nonlinear spatial transformation techniques. This can be explained by noting that the template brain has lost anatomical details due to extensive averaging. In order to bring functional data of a subject into Talairach space, the obtained spatial transformation for the anatomical data may be applied to the functional data if it has been coregistered with the unnormalized anatomical data set. Using the intensity-driven matching approach, functional data sets may also be directly normalized (without the help of anatomical data sets) because versions of the MNI template brain for functional (EPI) scans are also available. If possible, it is, however, recommended to apply the transformation obtained for the anatomical data also to the functional data because this approach guarantees that the precision of functional-anatomical alignment achieved during coregistration is not changed during the normalization step. More advanced volume-based normalization schemes have been proposed that replace the presented simple intensity-driven approaches (e.g., DARTEL, Ashburner 2007).

3.5.2 Cortex-Based Normalization

In recent years, more advanced brain normalization techniques have been proposed going beyond simple volume space alignment approaches. A particular interesting method attempts to explicitly align the cortical folding pattern (macroanatomy) across subjects (Fischl et al. 1999; Goebel et al. 2004, 2006; Frost and Goebel 2012) starting with topologically correct cortex mesh representations (see Sect. 3.4.3). The folded cortex meshes are first morphed to spherical representations since the restricted space of a sphere allows alignment using only two dimensions (longitude and latitude) instead of three dimensions as needed in volume space. Since the inflation of cortex hemispheres to spheres removes ("flattens") information of the gyral/sulcal folds, the respective information is retained by calculating curvature maps prior to inflation that are projected on the spherical representations. Cortex meshes from different subjects are then aligned on the sphere by *increasing the overlap of curvature information*. Since the curvature of the cortex reflects the gyral/sulcal folding pattern of the brain, this brain matching approach essentially aligns corresponding gyri and sulci across brains. It has been shown that cortex-based alignment substantially increases the statistical power and spatial specificity of group analyses by increasing not only the overlap of macroanatomical regions but also the overlap of corresponding functionally defined specialized brain areas (Frost and Goebel 2012).

3.5.3 Correspondence Based on Functional Localizer Experiments

An interesting approach to establish correspondence between brains is to use functional information directly. Using standardized stimuli, a specific region-of-interest (ROI) may be functionally identified in each subject. The ROIs identified in such *functional localizer experiments* are then used to extract time courses in subsequent main experiments. The extracted time courses of individual subjects are then integrated in group analyses (see below). If the assumption is correct that localizer experiments reveal corresponding brain regions in different subjects, the approach provides an optimal solution to the correspondence problem and will allow detection of subtle differences in fMRI responses at the group level with high statistical power. Statistical sensitivity is further enhanced by avoiding the massive multiple comparison correction problem. Instead of hundreds of thousand voxel-wise tests, only a few tests have to be performed – one for each considered ROI. The approach is statistically sound (no circularity) because the considered regions have been determined *independently* from the main data using special localizer runs. It may also be acceptable to use the same functional data for both localizer and main analysis as long as the contrast to localize ROIs is orthogonal to any contrast used to statistically test more subtle differences. The localizer approach has been applied successfully in many experiments, most notably in studies of the ventral visual cortex (e.g., O’Craven and Kanwisher 2000).

Unfortunately, it is often difficult to define experiments localizing the same pattern of activated brain areas in all subjects, especially in studies of higher cognitive functions, such as attention, mental imagery, working memory, and planning. If at all possible, the selection of corresponding functional brain areas in these experiments is very difficult and depends on the investigator’s choice of thresholding statistical maps and often on additional decisions such as grouping subclusters to obtain the same number of major clusters for each subject. Note that the increased variability of activated regions in more complex experiments could be explained by at least two factors. On the one hand the location of functionally corresponding brain regions may vary substantially across subjects with respect to aligned macro-anatomical structures. On the other hand, subjects may engage in different cognitive strategies to solve the same task leading to a (partially) different set of activated brain areas. Most likely, the observed variability is caused by a mixture of both sources of variability. Another problem of the localizer approach is the tendency to focus only on a few brain areas, namely, those, which can be mapped consistently in different subjects. This tendency bears the danger to overlook other important brain regions. This can be avoided by a recently proposed approach, functionally informed cortex-based alignment (Frost and Goebel 2013), that integrates

ROI-based and whole-cortex analysis using a modified version of cortex-based alignment that uses corresponding pre-mapped ROIs as alignment targets in addition to anatomical curvature information.

3.5.4 Statistical Analysis of Group Data

After brain normalization, the whole-brain data from multiple subjects can be statistically analyzed simply by concatenating time courses at corresponding locations. The corresponding locations can be voxel coordinates in Talairach/MNI space, vertex coordinates in cortex space, or identified ROIs in the localizer approach. Note that the power of statistical analysis depends on the quality of brain normalization. If the achieved alignment of corresponding functional brain areas is poor, suboptimal group results may be obtained since active voxels of some subjects will be averaged with non-active voxels (or active voxels from a non-corresponding brain area) from other subjects. In order to increase the overlap of activated brain areas across subjects in volume space, the functional data of each subject is often smoothed, typically using rather large Gaussian kernels with a full width at half maximum (FWHM) of 8–12 mm. While such an extensive spatial smoothing increases the overlap of active regions, it introduces other problems including potential averaging of non-corresponding functional areas within and across brains; furthermore, functional clusters smaller than the smoothing kernel will be suppressed. While spatial smoothing may be beneficial to reduce noise, it may also reduce detection sensitivity of truly active but small functional clusters. Extensive spatial smoothing may not be necessary when using advanced volumetric normalization schemes (e.g., Ashburner 2007), cortex-based alignment (e.g., Frost and Goebel 2012, 2013), or functional localizers.

After concatenating the data, the statistical analysis described for single-subject data (see Sect. 3.3) can be applied to the integrated time courses. In the context of the GLM, the multi-subject voxel time courses as well as the multi-subject predictors may be obtained by concatenation. After estimating the beta values, contrasts can be tested in the same way as described for single-subject data. While the described concatenation approach leads to a high statistical power due to the large number of blocks or events, the obtained results cannot be generalized to the population level since the data is analyzed as if it stems from a single subject. Significant findings only indicate that the results are replicable for the same “subject” (group of subjects). In order to test whether the obtained results are valid at the population level, the statistical procedure must assess the variability of observed effects across subjects (*random effects analysis*) as opposed to the variability across individual measurement time points as performed in the concatenation approach (*fixed effects analysis*). There are many statistical methods to

assess the variability across subjects for the purpose of proper population inferences. A simple and elegant method is provided by *multilevel summary statistics approach* (e.g., Kirby 1993; Holmes and Friston 1998; Worsley et al. 2002; Beckmann et al. 2003; Friston et al. 2005). In the first analysis stage, parameters (*summary statistics*) are estimated for each subject independently (level 1, fixed effects). Instead of the full time courses, only the resulting first-level parameter estimates (betas) from each subject are carried forward to the second analysis stage where they serve as the *dependent variables*. The second-level analysis assesses the consistency of effects within or between groups based on the variability of the first-level estimates across subjects (level 2, random effects). This hierarchical analysis approach reduces the data for the second stage analysis enormously since the time course data of each subject has been “collapsed” to only one or a few parameter estimates per subject. Since the summarized data at the second level reflects the variability of the estimated parameters across subjects, obtained significant results can be generalized to the population from which the subjects were drawn as a random sample.

To summarize the data at the first level, a standard GLM may be used to estimate effects – beta values – separately for each subject. Instead of one set of beta values in fixed effects analysis, this step will provide a *separate set of beta values for each subject*. The obtained beta values can be analyzed at the second level using again a GLM or a standard ANOVA with one or more within-subjects factors categorizing the beta values. If the data represent multiple groups of subjects, a between factor for group comparisons can be added.

These short explanations indicate that the statistical analysis at the second level does not differ from the usual statistical approach in medical studies. The only major difference to standard statistics is that the analysis is performed separately for each voxel (or vertex) requiring correction for a massive multiple comparison problem as has been described above. Note that in addition to the estimated subject-specific effects of the fMRI design (beta values of first-level analysis), additional external variables (e.g., an IQ value for each subject) may be incorporated as covariates at the second level.

3.6 Selected Advanced Data Analysis Methods

The analysis steps described in previous sections for single subjects and for group comparisons represent essential components of a standard fMRI analysis, which are performed in a similar way for most fMRI studies. Such a standard analysis involves proper preprocessing that includes drift removal and 3D motion correction, coregistration of functional and anatomical data, brain normalization, and a thorough statistical analysis usually based on the general linear model. The

standard procedure produces statistical maps that localize regions showing differential responses with respect to specified experimental hypotheses. Random effects group analyses allow generalization of observed findings from a sample of subjects to the population level. Event-related averages of active brain regions or prespecified ROIs can be used to compare estimated condition time courses within and across brain areas, often revealing additional interesting insights. The following sections shortly describe a selected list of further analysis methods aiming at improving or extending the standard analysis procedure.

3.6.1 Nonparametric Statistical Approaches

As stated in Sect. 3.3.3.5, GLM hypothesis testing requires normally distributed residuals with equal variance. Fortunately, the GLM is robust with respect to minor violations of the normality assumption. To avoid, however, wrong inferences due to non-normal distributions, nonparametric methods may be used, especially when analyzing small data samples.

3.6.2 Bayesian Statistics

It has been proposed to use Bayesian statistics because it provides an elegant framework for multilevel analyses (Friston et al. 2002). In the Bayesian approach, the data of a single experiment (or the data of a single subject) is not considered in isolation, but in light of available a priori knowledge. This a priori knowledge is formalized with prior probabilities $p(H_i)$ for relevant initial hypotheses H_i . Obtained new data D modifies the a priori knowledge resulting in posterior conditional probabilities $p(H_i|D)$, which are updated probabilities of the initial hypotheses given the new data. To calculate these probabilities, the inverse conditional probabilities $p(D|H_i)$ must be known describing the probability to obtain certain observations given that the hypotheses H_i are true. In the context of the *empirical Bayes approach*, these conditional probabilities can be estimated from the data. The empirical Bayes approach is appropriate for the analysis of fMRI data, since it allows an elegant formulation of hierarchical random effects analyses. It is, for example, possible to enter estimated parameters at a lower level as prior probabilities at the next higher level. Furthermore, the approach allows integrating correction for multiple comparisons resulting in thresholding values similar to the ones obtained with the false discovery rate approach.

3.6.3 Brain Normalization

As described earlier, brain normalization methods have an important influence on the quality of group analyses, since optimization of the standard analysis does not lead to substantial improvements if voxel time courses are concatenated from nonmatching brain regions. The described cortex-based normalization technique may substantially improve the alignment of homologue brain regions across

subjects. For more complex tasks, different, nonmatching activity patterns might reflect different cognitive strategies used by subjects. To cope with this situation, it would be desirable to use methods allowing automatic estimation of the similarity of activity patterns across subjects. Such methods could suggest splitting a group in subgroups with different statistical maps corresponding to the neural correlate of different cognitive strategies. Such a clustering approach has been recently implemented in the context of group-level ICA analyses (Esposito et al. 2005).

3.6.4 Data-Driven Analysis Methods

When considering the richness of fMRI data, it may be useful to apply *data-driven analysis methods*, which aim at discovering interesting spatiotemporal relationships in the data, which would be eventually overlooked with a purely hypothesis-driven approach. Data-driven methods, such as *independent component analysis* (ICA, e.g., McKeown et al. 1998a, b; Formisano et al. 2002), do not require a specification of expected, stimulus-related responses since they are able to extract interesting information automatically (“blindly”) from the data. It is, thus, not necessary to specify an explicit statistical model (design matrix). This is particularly interesting with respect to paradigms for which the exact specification of event onsets is difficult or impossible. Spatial ICA of fMRI data has been successfully applied in many tasks including the automatic detection of active networks during perceptual switches of ambiguous stimuli (Castelo-Branco et al. 2002) and the automatic detection of spontaneous hallucinatory episodes in schizophrenic patients (van de Ven et al. 2005). Data-driven methods are exploratory in nature and should not be viewed as replacements but as complementary tools for hypothesis-driven methods: If interesting, unexpected events have been discovered with a data-driven method, one should test these observations in succeeding studies with a hypothesis-driven standard statistical analysis. More generally, ICA has become an important method to reveal functionally connected networks, especially in the context of resting-state fMRI (see below).

3.6.5 Multivariate Analysis of Distributed Activity Patterns

Multi-voxel pattern analysis (MVPA) is gaining increasing interest in the neuroimaging community because it allows to detect differences between conditions with higher sensitivity than conventional univariate analysis by focusing on the analysis and comparison of distributed patterns of activity (Haxby et al. 2001). In such a multivariate approach, data from individual voxels within a region are jointly analyzed. Furthermore, MVPA is often presented in the context of “brain reading” applications reporting that specific mental

states or representational content can be decoded from fMRI activity patterns after performing a “training” or “learning phase.” In this context, MVPA tools are often referred to as classifiers or, more generally, learning machines. The latter names stress that many MVPA tools originate from a field called machine learning, a branch of artificial intelligence. In fMRI research, the support vector machine (SVM, Vapnik 1995) has become a particular popular machine learning classifier, which is used both for analyzing patterns in ROIs and for discriminating patterns that are potentially spread out across the whole brain.

Another popular MVPA approach is the “searchlight” method (Kriegeskorte et al. 2006). In this approach, each voxel is visited, as in a standard univariate analysis, but instead of using data of the visited voxel only for analysis, several voxels in the neighborhood are included forming a set of features for joined multivariate analysis. The neighborhood is usually defined roughly as a sphere, that is, voxels within a certain (Euclidean) distance from the visited voxel are included. The result of the multivariate analysis is then stored at the visited voxel (e.g., a *t* value resulting from a *multivariate* statistical comparison or an accuracy value from a support vector machine classifier). By visiting all voxels and analyzing their respective (partially overlapping) neighborhoods, one obtains a whole-brain map in the same way as when running univariate statistics.

3.6.6 Real-Time Analysis of fMRI Data

The described steps and techniques to analyze functional MRI data are very computation intensive and are, thus, performed in most cases hours or days after data acquisition has been completed. There are many scenarios that would benefit greatly from a *real-time analysis of fMRI data*, especially when studying single subjects as in presurgical mapping. Using appropriately modified analysis tools and state-of-the-art computer hardware, it is nowadays possible to perform real-time fMRI analysis during an ongoing experiment, including 3D motion correction and incremental GLM statistics of whole-brain recordings (Goebel 2012; Weiskopf 2012). It is even possible to run multivariate data-driven tools in real-time, including ICA (Esposito et al. 2003) and multi-voxel pattern analyses (LaConte et al. 2007; Sorger et al. 2010). One obvious benefit of real-time fMRI analysis is *quality assurance*. If, for example, one observes during an ongoing measurement that a patient moves too much or that the (absence of) activity patterns indicates that the task was not correctly understood, the running measurement may be stopped and repeated after giving the subject further instructions. If the ongoing statistical analysis on the other hand indicates that expected effects have reached a desired significance level earlier than expected, one could save scanning time by stopping the

measurement ahead of schedule. Real-time fMRI offers also the possibility to plan optimal slice positioning for subsequent runs based on the results obtained of an initial run. Based on the results of a first run, it would be, for example, possible to position a small slab of slices at an identified functional region for subsequent high-resolution spatial and/or temporal scanning. More advanced applications of real-time fMRI include neurofeedback (Weiskopf et al. 2003) and communication BCIs (Sorger et al. 2012). In fMRI neurofeedback studies, subjects learn to voluntarily control the level of activity in circumscribed brain areas by engaging in mental tasks such as inner speech, visual or auditory imagery, spatial navigation, mental calculation, or recalling (emotional) memories. In recent years, fMRI neurofeedback has been successfully employed as a therapeutic tool for various psychiatric and neurological diseases (e.g., Linden et al. 2012; Subramanian et al. 2011).

4 Functional Connectivity and Resting-State Networks

Generally, three types of brain connectivity are distinguished in brain research (Sporns 2010). *Structural connectivity* (or anatomical connectivity) refers to the physical presence of an axonal projection from one brain area to another. This type of connectivity and how diffusion MRI and computational tractography can be used to identify large axon bundles in the human brain is described in Sect. 5. *Functional connectivity* refers to the correlation structure in the data that can be used to reveal functional coupling between specific brain regions and to reveal functional networks. Finally, *effective connectivity* refers to models that go beyond correlation (or more generally statistical dependency) to more advanced measures of directed influence and causality within networks (Friston et al. 1994).

4.1 Functional and Effective Connectivity

Functional and effective connectivity methods aim to reveal the *functional integration* of brain areas, whereas the classical voxel-wise statistical approach (Sect. 3) is suited to reveal the *functional segregation* (functional specialization) of brain regions. Besides data-driven methods such as independent component analysis (ICA), many approaches have been used to model the interaction between spatially remote brain regions more explicitly. In the simplest case, the time courses from two regions are correlated resulting in a measure (e.g., linear correlation coefficient) of functional connectivity. Functional connectivity can be calculated separately for

different experimental conditions, which allows to assess whether two brain areas change their *functional coupling* in different cognitive contexts (Büchel et al. 1999). In conditions of attention, for example, two remote areas might work more closely with each other than in conditions of no attention.

Models of effective connectivity go beyond simple pairwise correlation analysis and assess the validity of models containing directed interactions between brain areas. These directed effective connections are often symbolized by arrows connecting boxes each representing a different brain area. Structural equation models (SEM, e.g., McIntosh and Gonzalez-Lima 1994) and, more recently, dynamic causal modeling (DCM, e.g., Penny et al. 2004) are used to test effective connectivity models. An interesting *data-driven* approach to effective connectivity modeling is provided by methods based on the concept of *Granger causality*. This approach does not require specification of connectivity models but enables to automatically detect effective connections from the data by mapping Granger causality for any selected reference voxel or region-of-interest (Goebel et al. 2003; Roebroeck et al. 2005, 2011).

4.2 Resting-State Networks

In recent years, functional connectivity studies have gained increased interest where the subject is in a relaxed resting state, that is, in the absence of experimental tasks and behavioral responses. These *resting-state fMRI* (RS-fMRI) studies allow measuring the amount of spontaneous BOLD signal synchronization within and between multiple regions across the entire brain (Biswal et al. 1995). The measured RS-fMRI activity is characterized by low-frequency (0.01–0.1 Hz) BOLD signal fluctuations, which are topologically organized as multiple spatially distributed functional connectivity networks called *resting-state networks* (RSNs) (e.g., van de Ven et al. 2004; De Luca et al. 2006). Spatial ICA (see Sect. 3.6.4) at the individual and group level is commonly applied in resting-state fMRI revealing RSNs that are consistently found in individuals, including the default-mode network (often separated in an anterior and posterior subnetwork), a visual and a auditory network, a sensorimotor network, and two (lateralized) dorsolateral frontoparietal networks (Fig. 23; for further details, see, e.g., Allen et al. 2011). The extracted independent components are usually scaled to spatial *z*-scores (i.e., the number of standard deviations of their whole-brain spatial distribution). These values express the relative amount a given voxel is modulated by the activation of the component (McKeown et al. 1998b) and hence reflect the amplitude of the correlated fluctuations within the corresponding functional connectivity network.

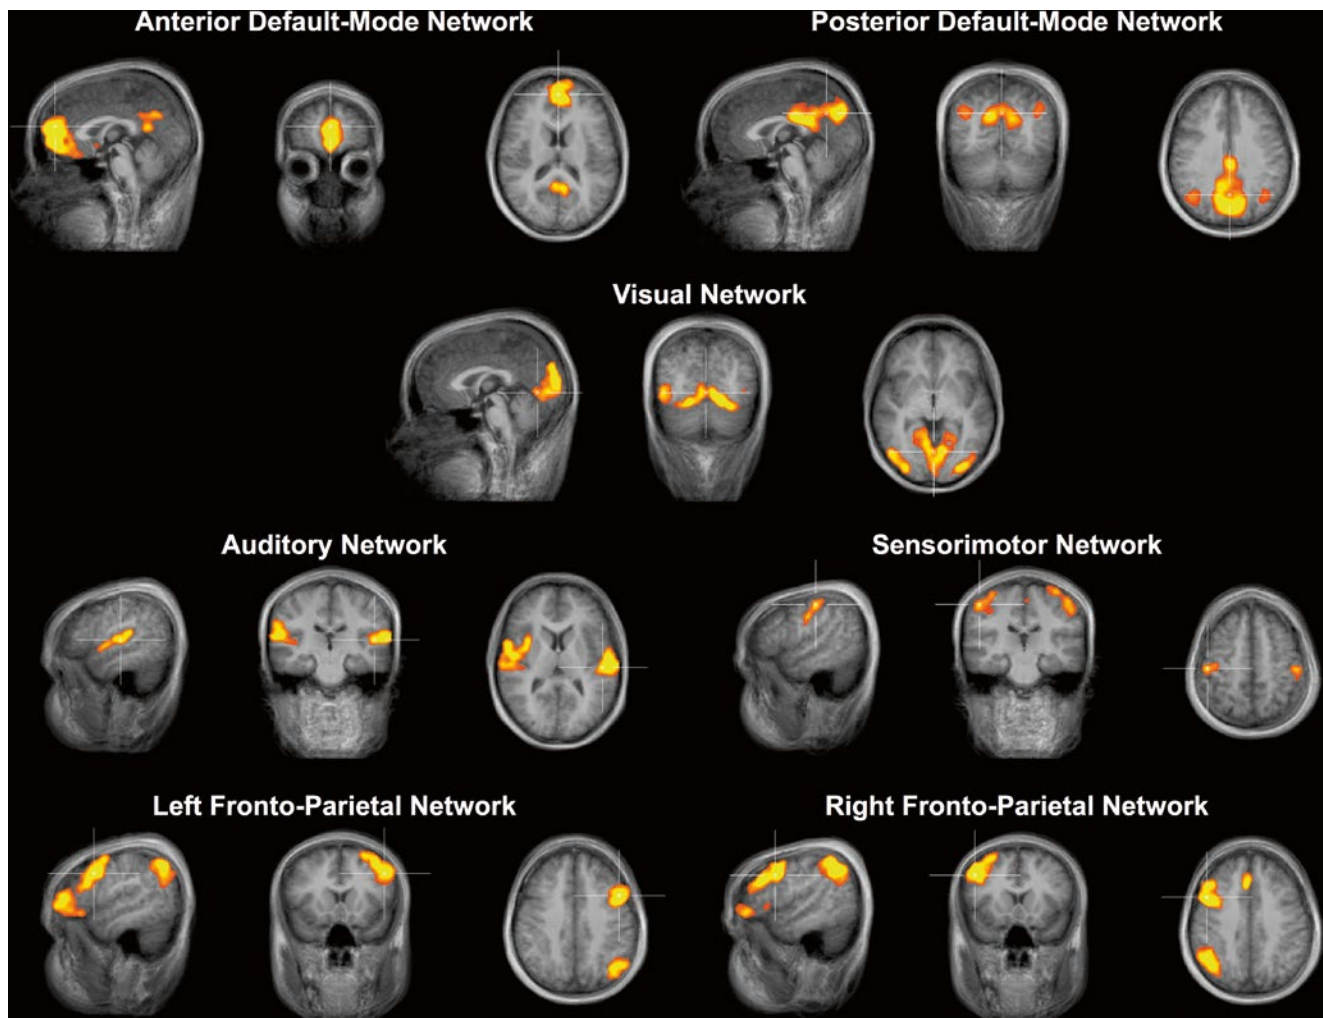


Fig. 23 A subset of major resting-state networks (RSNs) obtained by ICA analysis of the resting-state fMRI data of a group of healthy individuals ($n=8$); the default-mode network (DMN) is split in an anterior and posterior part (*upper row*)

An alternative (less objective) approach to retrieve RSNs is to calculate whole-brain correlations from seed regions that correspond to core locations of RSNs. In this approach the DMN, for example, can be retrieved by selecting a region in the posterior cingulate cortex as the seed region and then correlating each voxel's time course with the reference time course from the seed region. For both the seed-based correlation and ICA approach, it is recommended to account for possible BOLD effects due to cardiac pulsation and respiratory cycle (Birn et al. 2008) using a physiological noise correction method such as RETROICOR (Glover et al. 2000).

The obtained functional networks during rest conditions demonstrate that the brain is never “at rest” and the description of RSNs is, thus, a useful approach to explore the brain's functional organization in healthy individuals as well as to examine if it is altered in neurological or psychiatric diseases. Furthermore, it has been possible to relate RSNs to externally modifiable factors, such as different pharmacological treat-

ments or psychological experiences (Khalili-Mahani et al. 2012; Esposito et al. 2014). The default-mode network (DMN) has gained particular attention – the term “default mode” has been introduced by Raichle et al. (2001) to describe resting-state brain function. The DMN is a network of brain regions that include part of the medial temporal lobe (presumed memory functions), part of the medial prefrontal cortex (presumed theory of mind functions), the posterior cingulate cortex along with the adjacent ventral precuneus, and the medial, lateral, and inferior parietal cortex. The DMN is active when the individual is not focused on the outside world and the brain is at wakeful rest corresponding likely to task-independent introspection, mind-wandering, and self-referential thought. During goal-oriented activity, the DMN is deactivated and other regions are active that are sometimes described as the task-positive network (TPN). The DMN has been hypothesized to be relevant to disorders including Alzheimer's disease, autism, and schizophrenia (Buckner et al. 2008).

5 Diffusion-Weighted MRI and Tractography

In recent years, MRI has not only revolutionized functional brain imaging targeting gray matter neuronal activity but also enabled insights into the human white matter structure using diffusion-weighted magnetic resonance imaging (DW-MRI, dMRI, or DWI). Pulse sequences for dMRI measure the diffusion of water molecules in each voxel providing information about the fibers in that voxel that can be used to assess the “intactness” of the white matter structure and as the basis for computational tractography since the diffusion process is hindered by the boundaries of the fibers forcing the majority of water molecules to diffuse along these fibers.

A diffusion-weighted MR measurement consists of several volumes each measuring the reduction of the signal resulting from diffusion along a specific axis in space that is selected by setting the x , y , and z gradients of the scanner accordingly using a pulsed-gradient spin echo-sequence (PGSE) developed by Stejskal and Tanner (1965).

5.1 Diffusion Tensor Imaging

It has been proposed to model the diffusion measured in a voxel as a 3D Gaussian probability function from which a *diffusion tensor* (3×3 matrix) can be calculated (Basser et al. 1994), which has led to the name *diffusion tensor imaging* (DTI) for the most widely used diffusion-weighted MRI acquisition and modeling approach. In order to construct the diffusion tensor, a minimum of six diffusion-weighted volumes and a non-diffusion-weighted image need to be measured. From the diffusion tensor, the principal diffusion directions (three eigenvectors of the tensor) and associated diffusion coefficients (three eigenvalues λ_1 , λ_2 , λ_3) can be derived. Note that although eigenvectors mathematically represent directions, DTI cannot distinguish opposing directions from each other, that is, the resulting values estimate diffusion along opposing directions, that is, along principle axes of diffusion. The eigenvectors and eigenvalues can be visualized as an ellipsoid. If water molecules diffuse without restrictions in all directions, the resulting “ellipsoid” will have the shape of a sphere, that is, all three axes (eigenvectors) have the same length ($\lambda_1 = \lambda_2 = \lambda_3$) and there is no preferred axis of diffusion. This situation is described as *isotropic diffusion*. In case that water molecules diffuse with low restrictions along one axis but diffusion is hindered in other directions, a strongly elongated (cigar shaped) ellipsoid will be obtained ($\lambda_1 \gg \lambda_2 \approx \lambda_3$). This case of restricted diffusion occurs within and around white matter fibers and is described as *anisotropic diffusion*. In this case, the main (longest) axis of the resulting ellipsoid will likely coincide

with the main orientation of fiber bundles running through the measured voxel. This is the principle assumption of DTI. Note, however, that the tensors estimated in each voxel do not provide fibers but only local discrete measurements, that is, putative fibers need to be reconstructed using *computational tractography*, that is, the orientation of estimated tensors need to be “concatenated” across neighboring voxels. Since results of specific tractography procedures are dependent on many factors (see below), visualized fibers need to be interpreted with care.

Several interesting quantities can be derived from the diffusion tensor in each voxel. The *mean diffusivity* quantifies the overall movement of water molecules in a voxel, which depends on tissue type (e.g., CSF vs. the white matter) and the presence of diffusion restrictions (e.g., axons). Figure 2.24b shows that the mean diffusivity is high in the ventricles (yellow color) while it is low in the white and gray matter (orange color). The most common derived scalar quantity is *fractional anisotropy* (FA) that characterizes the overall shape of the diffusion, that is, it quantifies the fraction of the diffusion tensor that can be ascribed to anisotropic diffusion:


$$FA = \frac{\sqrt{(\lambda_1 - \lambda_2)^2 + (\lambda_2 - \lambda_3)^2 + (\lambda_1 - \lambda_3)^2}}{\sqrt{2(\lambda_1^2 + \lambda_2^2 + \lambda_3^2)}}$$

The FA value varies between 0 (isotropic diffusion, shape of a sphere) and 1 (maximal anisotropy, shape of a line). Figure 2.24c shows that fractional anisotropy is high (yellow color) in the white matter (e.g., in the corpus callosum) but low (orange color) in the gray matter and ventricles. The FA value disregards the specific diffusion axis. A value of 0 indicates no preferred diffusion axis (sphere) while a value of 1 indicates diffusion precisely along a single axis. Since the white matter contains parallel fibers within larger tracts, it contains usually high FA values (>0.3) whereas FA values are low (0.0–0.2) in the gray matter. The FA quantity has gained increasing interest in recent years since it has been shown that FA values in specific tracts can be related to specific diseases and since they correlate with cognitive performance measures such as reading capability (see Sect. 5.3).

5.1.1 Tractography: From Tensors to Fiber Bundles

Based on the preferred orientation of the tensors in neighboring voxels, computational tractography or *fiber tracking* procedures aim to reconstruct the trajectory of fibers in the white matter by “concatenating” neighboring tensors. Fiber tracking is usually launched (seeded) in all voxels (even in sub-voxel coordinate grids) except those with low FA values since they do not reflect strong directness. The tracking process then generates a large amount of short and long reconstructed (“soft-

Fig. 2.24 Important voxel-wise measures that can be extracted from diffusion-weighted MRI scans. (a) Anatomical scan shown as reference. (b) Mean diffusivity map coregistered with anatomy shown in (a); note that diffusivity is high in CSF (ventricles, *yellow color*) but low in *gray matter* and *white matter* fiber bundles such as the corpus callosum (*orange color*). (c) Fractional anisotropy (FA) map coregistered with anatomy shown in (a); note that FA is low (*orange color*) in the presence of low diffusion restrictions (ventricles) but high in white matter fiber bundles such as the corpus callosum containing coherently oriented fibers within voxels



ware”) fibers. Specific fiber tracts are extracted from the dense fiber field by using regional constraints (e.g., Catani and Thiebaut de Schotten 2008), that is, fibers belonging to a specific tract are included if they pass through one or more specified volumes-of-interest (VOIs). Since the main axis of the tensor indicates an oriented axis and not a direction, fiber tracking is performed in two opposing directions. After both “half-fibers” have been reconstructed, they are finally integrated into a single fiber. In order to reconstruct a (half-) fiber, a small (sub-voxel) step is performed in one of the two directions provided by the main (longest) axis of the ellipsoid at a seed position. At the reached position, the direction for the next small step will be calculated using the tensor orientation and the direction of the previous step. Since the reached position usually does not correspond to integral coordinates (i.e., it falls between voxels), the calculation of the next direction is based on the tensors surrounding the current 3D position; in this interpolation process, tensors influence the calculation with respect to the distance of the corresponding voxels to the current position. After updating the direction, the next step is performed. Again a new direction is calculated at the new position for the next step and so on producing a connected trajectory of short line segments. This process continues until certain stop criteria are reached such as when an FA value is reached that falls below a specified threshold or in case that the reconstructed fiber leaves the white matter. In order to create smooth reconstructed fibers (Fig. 25), the chosen step size needs to be smaller than the distance of the voxels.

5.2 Validation and Improvements

While tractography usually creates interesting results, it is important to realize that visualized fibers are reconstructed from diffusion estimates that are measured at discrete 3D positions (voxels) and, thus, may not necessarily reflect true fiber tracts in the brain. A central concern in current tractography research concerns the question how much one can trust the beautiful pictures generated by fiber tracking procedures. The answer depends on many factors including the quality of the diffusion-weighted measurement which is influenced by scanner parameters (e.g., signal-to-noise ratio) as well as by parameters of the participant such as head motion and physiological noise. The most important limiting factor is related to the voxel size used for in vivo studies that is a few orders of magnitudes larger than the small scale at which the diffusion of water molecules happens. With a typical spatial resolution of about 2 mm, only the average diffusion of water molecules in a large cube (voxel) is captured, which does not allow to resolve fine-grained white matter fiber bundles or fiber bundles in the gray matter. The resolution issue relates also to the “kissing or crossing” problem, that is, it can often not be decided in a large voxel whether two (or more) incoming fiber bundles cross in that voxel or whether they merely touch each other and part by changing direction.

Despite its usefulness in many applications, the diffusion tensor model has the drawback of being a unidirectional model. Its orientation estimation works very well in areas characterized by prominent fiber pathways following one direction, giving rise to a unimodal water diffusivity profile. When, however, several different diffusion directions are present in one voxel, the estimated diffusion tensor contains directionality information which has high uncertainty at best (low precision) or is even biased to a wrong average orientation. In order to obtain more valid results from diffusion-weighted measurements, several advanced measurement schemes and analysis methods have been proposed. The most complete approach to estimate the full fiber orientation density function is diffusion spectrum imaging (DSI) that requires, however, very long measurement times (Wedeen et al. 2005). Somewhat less time-consuming advanced approaches are q-ball imaging (Tuch 2004) and spherical deconvolution (Tournier et al. 2004). These modeling approaches go beyond the simple tensor model and fit more complex models to the measured diffusion data that no longer assume a single major diffusion axis but explicitly allow multiple (crossing) fibers in a voxel. In order to provide sufficient constraints for these more complex models, many more diffusion directions (e.g., 100) need to be measured as for conventional DTI scans that require only 6 diffusion directions. Because of the high number of direction measurements, these approaches are also called “HARDI” (high angular resolution diffusion imaging) methods. Since HARDI measurements (Tuch 2002) need much longer scanning time than DTI measurements, they are not common in clinical MRI measurements. Even with more advanced measurement and analysis approaches, reconstructed fiber tracts may vary substantially depending on the used tractography algorithm (Bastiani et al. 2012). To validate fiber tracking algorithms, it is important to have ground-truth data, that is, knowledge about the true trajectory of fiber bundles. One way to perform ground-truth validation is to use “DTI phantoms” that contain known, artificially created fibers, including challenging cases with crossing and kissing fibers (Pullens et al. 2010). Another important validation approach uses postmortem brain tissue that is analyzed both with dMRI as well as with tracers that are released in specific brain areas. Since these tracers traverse backward along axons through other regions, they reveal true region-to-region connectivity that can be used as ground-truth data for DWI-based connectivity analyses of the same tissue (e.g., Seehaus et al. 2013).

5.3 Applications

In recent years, diffusion MRI has led to several interesting applications. Especially the fractional anisotropy measure has become an important biomarker of white matter integrity serving as a local index to diagnose neurological or psychiatric diseases or to predict (lack of) cognitive performance. It

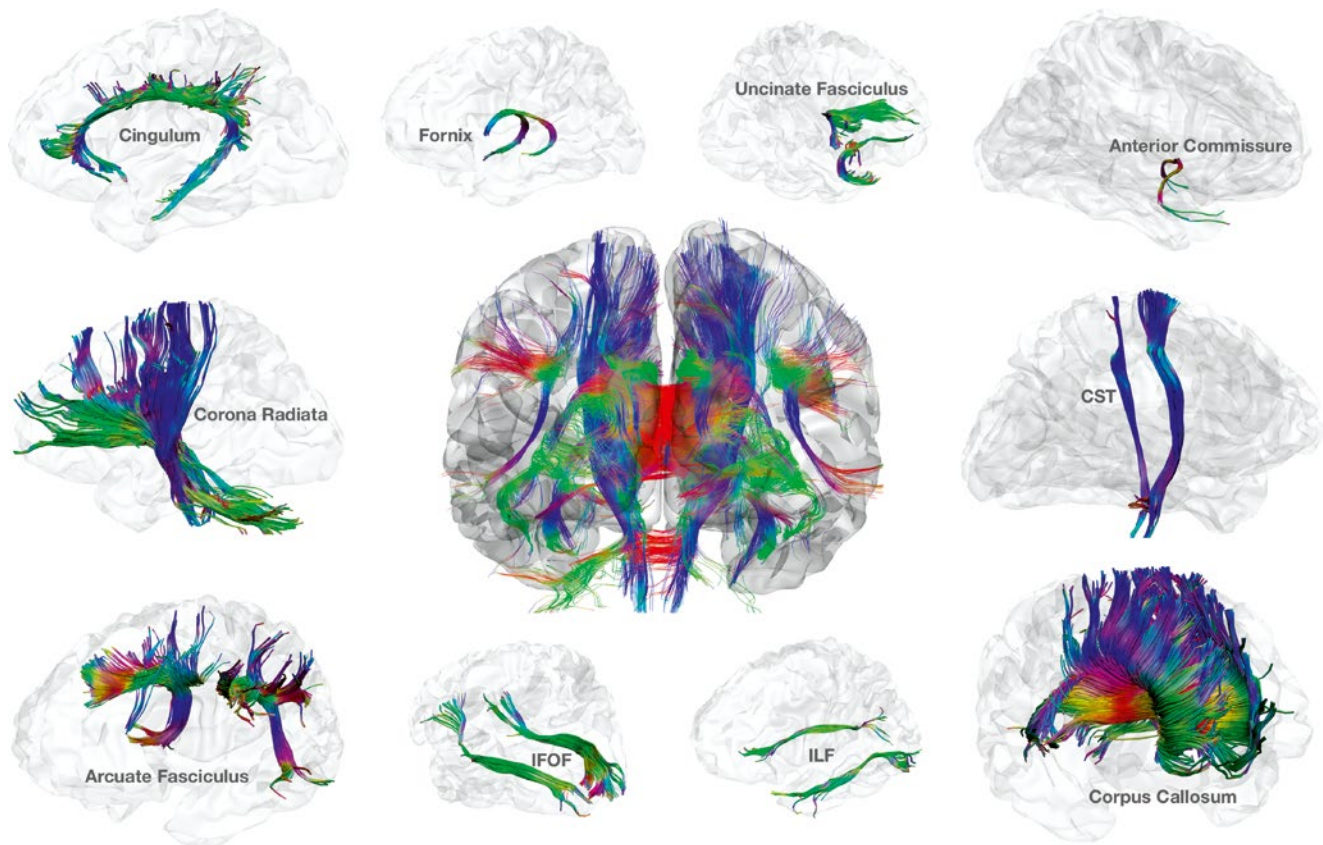


Fig. 25 A subset of major fiber tracts revealed by computational tractography from the diffusion-weighted MRI data of a healthy individual. *CST* corticospinal tract, *IFOF* inferior fronto-occipital fasciculus, *ILF* inferior longitudinal fasciculus

has been, for example, shown that FA values extracted from dMRI measurements from good and poor readers differ, and the size of the difference is largest within a region within the left hemisphere temporoparietal white matter (Klingberg et al. 2000; Deutsch et al. 2005). FA values could also be related to the level of creativity in several brain areas including the prefrontal cortex, basal ganglia, and at the border of the temporal and parietal lobe (Takeuchi et al. 2010). Note, however, that FA values are not fixed (“hard-wired”) properties of the white matter but can change depending on usage of the underlying fibers. It has been, for example, shown that FA values in the white matter in regions of the posterior parietal cortex (containing fibers that presumably mediate visuospatial transformation) significantly increase when subjects train on an intensive visual motor coordination task such as learning to juggle (Scholz et al. 2009). It has also been discovered that FA values reflect the development of cognitive abilities including systematic increases in the corpus callosum and prefrontal cortex during childhood (Barnea-Goraly et al. 2005); the changes observed in prefrontal cortical areas are discussed as related to the development of working memory, attention, and behavioral control. FA measures are also increasingly used for early diagnosis of stroke since reduced diffusion in

affected brain regions is often detected already minutes after the stroke. It is important to note that FA measurements are quantitative values (as opposed to fMRI measurements) that can be compared across people, labs, and scanners.

While computational tractography produces less objective results than FA estimates, reconstructed white matter fiber tracts are especially important to guide neurosurgical procedures potentially reducing the risk of lesioning important fiber tracts, for example, related to language functions. For this and similar purposes, several tools (e.g., Yeatman et al. 2012) are now available that allow extracting major long-range fiber tracts from dMRI data, including *commissural tracts* (e.g., corpus callosum) connecting both cortical hemispheres, *association tracts* (e.g., arcuate fasciculus) connecting regions within the same hemisphere, and *projection tracts* (e.g., corticospinal tract) connecting cortical regions to subcortical areas, the cerebellum, and the spinal cord. Figure 25 shows selected major fiber tracts that have been reconstructed from the dMRI data of a healthy individual; further details about the depicted (as well as other) fiber tracts are described, for example, by Catani and Thiebaut de Schotten (2008) and Yeatman et al. (2012).

5.3.1 The Human Connectome

An important aim of recent brain research is to understand how brain areas communicate with each other. This aim is pursued by investigating anatomical connectivity with dMRI to reconstruct in vivo the macroscale human connectome (Sporns et al. 2005), which is the map of all the structural connections in the human brain. This is complemented by functional connectivity studies using fMRI (see Sect. 4.1) and other modalities such as EEG and MEG. In integrative multimodal modeling approaches, the anatomical connectome may serve as an important structural constraint for functional connectivity models since only brain areas that are connected via fiber bundles may communicate directly with each other. Diffusion MRI may even help to estimate the strength of connectivity between brain areas. The currently most prominent attempt along these lines is the *Human Connectome Project* (<http://www.neuroscienceblueprint.nih.gov/connectome/>). This project aims to derive a complete map of all major connections between brain areas by measuring dMRI as well as functional connectivity and genetic data in more than 1,000 individuals (twin pairs and their siblings from 300 families). Besides deriving a connectivity map – the human connectome – the measured data of structural and functional connectivity will be shared to stimulate research in the emerging field of *human connectomics* as well as providing the basis for future studies of abnormal brain circuits in neurological and psychiatric disorders.

References

- Allen EA, Erhardt EB, Damaraju E, Gruner W, Segall JM et al (2011) A baseline for the multivariate comparison of resting-state networks. *Front Syst Neurosci* 5:2
- Ashburner J (2007) A fast diffeomorphic image registration algorithm. *Neuroimage* 38:95–113
- Ashburner J, Friston KJ (1999) Nonlinear spatial normalization using basis functions. *Hum Brain Mapp* 7:254–266
- Bandettini PA, Cox RW (2000) Event-related fMRI contrast when using constant interstimulus interval: theory and experiment. *Magn Reson Med* 43:540–548
- Bandettini PA, Birn RM, Donahue KM (2000) Functional MRI: background, methodology, limits, and implementation. In: Cacioppo JT, Tassinary LG, Berntson GG (eds) *Handbook of psychophysiology*. Cambridge University Press, New York, pp 978–1014
- Barnea-Goraly N, Menon V, Eckert M, Tamm L, Bammer R, Karchemskiy A, Dant CC, Reiss AL (2005) White matter development during childhood and adolescence: a cross-sectional diffusion tensor imaging study. *Cereb Cortex* 15:1848–1854
- Basser PJ, Mattiello J, LeBihan D (1994) Estimation of the effective self-diffusion tensor from the NMR spin echo. *J Magn Reson B* 103:247–254
- Bastiani M, Shah NJ, Goebel R, Roebroeck A (2012) Human cortical connectome reconstruction from diffusion weighted MRI: the effect of tractography algorithm. *Neuroimage* 62:1732–1749
- Beckmann CF, Jenkinson M, Smith SM (2003) General multilevel linear modeling for group analysis in FMRI. *Neuroimage* 20:1052–1063
- Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J Roy Statist Soc Ser B* 57:289–300
- Birn RM, Murphy K, Bandettini PA (2008) The effect of respiration variations on independent component analysis results of resting state functional connectivity. *Hum Brain Mapp* 29:740–750
- Biswal B, Yetkin FZ, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med* 34:537–541
- Blamire AM, Ogawa S, Ugurbil K et al (1992) Dynamic mapping of the human visual cortex by high-speed magnetic resonance imaging. *Proc Natl Acad Sci U S A* 89:11069–11073
- Bledowski C, Cohen Kadosh K, Wibrall M, Rahm B, Bittner RA, Hoehstetter K, Scherg M, Maurer K, Goebel R, Linden DE (2006) Mental chronometry of working memory retrieval: a combined functional magnetic resonance imaging and event-related potentials approach. *J Neurosci* 26:821–829
- 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
- Brown MA, Semelka RC (1999) *MRI – basic principles and applications*. Wiley-Liss, New York
- Büchel C, Coull JT, Friston KJ (1999) The predictive value of changes in effective connectivity for human learning. *Science* 283:1538–1541
- 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 U S A* 93:14878–14883
- Buckner RL, Andrews-Hanna JR, Schacter DL (2008) The brain’s default network: anatomy, function, and relevance to disease. *Ann N Y Acad Sci* 1124:1–38
- Bullmore E, Brammer M, Williams SC et al (1996) Statistical methods of estimation and inference for functional MR image analysis. *Magn Reson Med* 35:261–277
- Buxton RB, Wong EC, Frank LR (1998) Dynamics of blood flow and oxygen metabolism during brain activation: the balloon model. *Magn Reson Med* 39:855–864
- Buxton RB, Uludag K, Dubowitz DJ, Liu TT (2004) Modeling the hemodynamic response to brain activation. *Neuroimage* 23:S220–S233
- Castelo-Branco M, Formisano E, Backes W, Zanella F, Neuenschwander S, Singer W, Goebel R (2002) Activity patterns in human motion-sensitive areas depend on the interpretation of global motion. *Proc Natl Acad Sci U S A* 99:13914–13919
- Catani M, Thiebaut de Schotten M (2008) A diffusion tensor imaging tractography atlas for virtual in vivo dissections. *Cortex* 44:1105–1132
- Cheng K, Waggoner AK, Tanaka K (2001) Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. *Neuron* 32:359–374
- Cochrane D, Orcutt GH (1949) Application of least squares regression to relationships containing autocorrelated error terms. *J Am Stat Assoc* 44:32–61
- Dale AM, Buckner RL (1997) Selective averaging of rapidly presented individual trials using fMRI. *Hum Brain Mapp* 5:329–340
- Dale AM, Halgren E (2001) Spatiotemporal mapping of brain activity by integration of multiple imaging modalities. *Curr Opin Neurobiol* 11:202–208
- De Luca M, Beckmann CF, De Stefano N, Matthews PM, Smith SM (2006) fMRI resting state networks define distinct modes of long-distance interactions in the human brain. *Neuroimage* 29:1359–1367
- De Martino F, Moerel M, van de Moortele PF, Ugurbil K, Goebel R, Yacoub E, Formisano E (2013) Spatial organization of frequency preference and selectivity in the human inferior colliculus. *Nat Commun* 4:1386
- Deutsch GK, Dougherty RF, Bammer R, Siok WT, Gabrieli JD, Wandell B (2005) Children’s reading performance is correlated with white matter structure measured by diffusion tensor imaging. *Cortex* 41:354–363
- Draper NR, Smith H (1998) *Applied regression analysis*, 3rd edn. Wiley, New York

- Duvernoy HM, Delon S, Vannson JL (1981) Cortical blood vessels of the human brain. *Brain Res Bull* 7:519–579
- Esposito F, Seifritz E, Formisano E, Morrone R, Scarabino T, Tedeschi G, Cirillo S, Goebel R, Di Salle F (2003) Real-time independent component analysis of fMRI time-series. *Neuroimage* 20:2209–2224
- Esposito F, Scarabino T, Hyvarinen A, Himberg J, Formisano E, Comani S, Tedeschi G, Goebel R, Seifritz E, Di Salle F (2005) Independent component analysis of fMRI group studies by self-organizing clustering. *Neuroimage* 25:193–205
- Esposito F, Otto T, Zijlstra FRH, Goebel R (2014) Spatially distributed effects of mental exhaustion on resting-state FMRI networks. *PLoS One* 9:e94222
- Evans AC, Collins DL, Mills SR, Brown ED, Kelly RL, Peters TM (1993) 3D statistical neuroanatomical models from 305 MRI volumes. In: *Proceedings of IEEE-Nuclear Science Symposium and Medical Imaging Conference*. pp 1813–1817
- Fischl B, Sereno MI, Tootell RBH, Dale AM (1999) High-resolution inter-subject averaging and a coordinate system for the cortical surface. *Hum Brain Mapp* 8:272–284
- Forman SD, Cohen JD, Fitzgerald M, Eddy WF, Mintun MA, Noll DC (1995) Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a clustersize threshold. *Magn Reson Med* 3(5):636–647
- Formisano E, Goebel R (2003) Tracking cognitive processes with functional MRI mental chronometry. *Curr Opin Neurobiol* 13:174–181
- Formisano E, Esposito F, Kriegeskorte N, Tedeschi G, Di Salle F, Goebel R (2002) Spatial independent component analysis of functional magnetic resonance imaging time-series: characterization of the cortical components. *Neurocomputing* 49:241–254
- Fox PT, Raichle ME, Mintun MA, Dence C (1988) Nonoxidative glucose consumption during focal physiological neural activity. *Science* 241:462–464
- Friston KJ, Henson RN (2006) Commentary on: divide and conquer; a defense of functional localisers. *Neuroimage* 30:1097–1099
- Friston KJ, Jezzard P, Turner R (1994) The analysis of functional MRI time-series. *Hum Brain Mapp* 1:153–171
- Friston KJ, Holmes AP, Worsley KJ, Poline JP, Frith CD, Frackowiak RSJ (1995) Statistical parametric maps in functional imaging: a general linear approach. *Hum Brain Mapp* 2:189–210
- Friston KJ, Fletcher P, Josephs O, Holmes A, Rugg MD, Turner R (1998) Event-related fMRI: characterizing differential responses. *Neuroimage* 7:30–40
- Friston KJ, Penny W, Phillips C, Kiebel S, Hinton G, Ashburner J (2002) Classical and Bayesian inference in neuroimaging: theory. *Neuroimage* 16:465–483
- Friston KJ, Stephan KE, Lund TE, Morcom A, Kiebel S (2005) Mixed-effects and fMRI studies. *Neuroimage* 24:244–252
- Friston KJ, Rotshtein P, Geng JJ, Sterzer P, Henson RN (2006) A critique of functional localisers. *Neuroimage* 30:1077–1087
- Frost M, Goebel R (2012) Measuring structural-functional correspondence: spatial variability of specialised brain regions after macro-anatomical alignment. *Neuroimage* 59:1369–1381
- Frost M, Goebel R (2013) Functionally informed cortex based alignment: an integrated approach for whole-cortex macro-anatomical and ROI-based functional alignment. *Neuroimage* 83:1002–1010
- Fujita I, Tanaka K, Ito M, Cheng K (1992) Columns for visual features of objects in monkey inferotemporal cortex. *Nature* 360:343–346
- Genovese CR, Lazar NA, Nichols T (2002) Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage* 15:870–878
- Glover GH, Li TQ, Ress D (2000) Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magn Reson Med* 44:162–167
- Goebel R (2012) *BrainVoyager – past, present, future*. *Neuroimage* 62:748–756
- Goebel R, Singer W (1999) Cortical surface-based statistical analysis of functional magnetic resonance imaging data. *Neuroimage* 9, S64
- Goebel R, Roebroeck A, Kim DS, Formisano E (2003) Investigating directed cortical interactions in time-resolved fMRI data using vector autoregressive modeling and Granger causality mapping. *Magn Reson Imaging* 21:1251–1261
- Goebel R, Hasson U, Lefi I, Malach R (2004) Statistical analyses across aligned cortical hemispheres reveal high-resolution population maps of human visual cortex. *Neuroimage* 22(Suppl 2)
- Goebel R, Esposito F, Formisano E (2006) Analysis of functional image analysis contest (FIAC) data with BrainVoyager QX: from single-subject to cortically aligned group general linear model analysis and self-organizing group independent component analysis. *Hum Brain Mapp* 27:392–401
- Goebel R, Zilverstand A, Sorger B (2010) Real-time fMRI-based brain–computer interfacing for neurofeedback therapy and compensation of lost motor functions. *Imaging Med* 2:407–414
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293:2425–2430
- Holmes AP, Friston KJ (1998) Generalisability, random effects & population inference, in: *Fourth International Conference on functional mapping of the human brain*. *Neuro Image* 7:S754
- Huettel SA, Song AW, McCarthy G (2004) Functional magnetic resonance imaging. *Sinauer Associates*, Sunderland
- Khalili-Mahani N, Zoethout RM, Beckmann CF, Baerends E, de Kam ML et al (2012) Effects of morphine and alcohol on functional brain connectivity during “resting state”: a placebo-controlled crossover study in healthy young men. *Hum Brain Mapp* 33:1003–1018
- Kim DS, Duong TQ, Kim SG (2000) High-resolution mapping of iso-orientation columns by fMRI. *Nat Neurosci* 3:164–169
- Kirby KN (1993) *Advanced data analysis with SYSTAT*. Van Nostrand Reinhold, New York
- Klingberg T, Hedehus M, Temple E, Salz T, Gabrieli JD, Moseley ME, Poldrack RA (2000) Microstructure of temporo-parietal white matter as a basis for reading ability: evidence from diffusion tensor magnetic resonance imaging. *Neuron* 25:493–500
- Kriegeskorte N, Goebel R (2001) An efficient algorithm for topologically correct segmentation of the cortical sheet in anatomical MR volumes. *Neuroimage* 14:329–346
- Kriegeskorte N, Goebel R, Bandettini P (2006) Information-based functional brain mapping. *Proc Natl Acad Sci U S A* 103:3863–3868
- Kruggel F, von Cramon DY (1999) Temporal properties of the hemodynamic response in functional MRI. *Hum Brain Mapp* 8:259–271
- Kutner MH, Nachtsheim CJ, Neter J, Li W (2005) *Applied linear statistical models*, 5th edn. McGraw-Hill, Boston
- LaConte SM, Peltier SJ, Hu XP (2007) Real-time fMRI using brain-state classification. *Hum Brain Mapp* 28:1033–1044
- Lenoski B, Baxter LC, Karam LJ, Maisog J, Debbins J (2008) On the performance of autocorrelation estimation algorithms for fMRI analysis. *IEEE J Sel Top Sign Process* 2:828–838
- Linden DEJ, Habes I, Johnston SJ, Linden S, Tatineni R, Subramanian L, Sorger B, Healy D, Goebel R (2012) Real-time self-regulation of emotion networks in patients with depression. *PLoS One* 7:e38115
- Logothetis NK, Wandell B (2004) Interpreting the BOLD signal. *Annu Rev Physiol* 66:735–769
- Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412:150–157
- Magistretti PJ, Pellerin L, Rothman DL, Shulman RG (1999) Energy on demand. *Science* 283:496–497
- Mathiesen C, Caesar K, Lauritzen M (2000) Temporal coupling between neuronal activity and blood flow in rat cerebellar cortex as indicated by field potential analysis. *J Physiol* 523:235–246

- Maus B, van Breukelen GJ, Goebel R, Berger MP (2010a) Optimization of blocked designs in fMRI studies. *Psychometrica* 75:373–390
- Maus B, van Breukelen GJ, Goebel R, Berger MP (2010b) Robustness of optimal design of fMRI experiments with application of a genetic algorithm. *Neuroimage* 49:2433–2443
- McIntosh AR, Gonzalez-Lima F (1994) Structural equation modeling and its application to network analysis in functional brain imaging. *Hum Brain Mapp* 2:2–22
- McKeown MJ, Makeig S, Brown GG, Jung TP, Kindermann SS, Bell AJ, Iragui V, Sejnowski TJ (1998a) Blind separation of functional magnetic resonance imaging (fMRI) data. *Hum Brain Mapp* 6: 368–372
- McKeown MJ, Makeig S, Brown GG, Jung TP, Kindermann SS, Bell AJ, Sejnowski TJ (1998b) 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
- Moeller S, Yacoub E, Auerbach E, Strupp JP, Harel N, Ugurbil K (2010) Multi-band multi-slice GE-EPI at 7 Tesla, with 16 fold acceleration using partial parallel imaging with application to high spatial and temporal whole brain fMRI. *Magn Reson Med* 63:1144–1153
- Mulert C, Jager L, Schmitt R, Bussfeld P, Pogarell O, Moller HJ, Juckel G, Hegerl U (2004) Integration of fMRI and simultaneous EEG: towards a comprehensive understanding of localization and time-course of brain activity in target detection. *Neuroimage* 22:83–94
- NessAiver M (1997) All you really need to know about MRI physics. *Simply Physics*, Baltimore
- Nichols TE, Brett M, Andersson J, Wager T, Poline JB (2005) Valid conjunction inference with the minimum statistic. *Neuroimage* 25:653–660
- O’Craven KM, Kanwisher N (2000) Mental imagery of faces and places activates corresponding stimulus-specific brain regions. *J Cogn Neurosci* 12:1013–1023
- Ogawa S, Lee TM, Kay AR, Tank DW (1990) Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc Natl Acad Sci U S A* 87:9868–9872
- Penny WD, Stephan KE, Mechelli A, Friston KJ (2004) Comparing dynamic causal models. *Neuroimage* 22:1157–1172
- Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P (1999) SENSE: sensitivity encoding for fast MRI. *Magn Reson Med* 42:952–962
- Pullens P, Roebroeck A, Goebel R (2010) Ground truth hardware phantoms for validation of diffusion-weighted MRI applications. *J Magn Reson Imaging* 32:482–488
- Raichle ME, MacLeod AM, Snyder AZ, Powers WJ, Gusnard DA, Shulman GL (2001) A default mode of brain function. *Proc Natl Acad Sci U S A* 98:676–682
- Robson MD, Dorosz JL, Gore JC (1998) Measurements of the temporal fMRI response of the human auditory cortex to trains of tones. *Neuroimage* 7:185–198
- Roebroeck A, Formisano E, Goebel R (2005) Mapping directed influence over the brain using Granger causality and fMRI. *Neuroimage* 25:230–242
- Roebroeck A, Formisano E, Goebel R (2011) The identification of interacting networks in the brain using fMRI: model selection, causality and deconvolution. *Neuroimage* 58:296–302
- Saxe R, Brett M, Kanwisher N (2006) Divide and conquer: a defense of functional localizers. *Neuroimage* 30:1088–1096 (discussion 1097–1089)
- Scherg M, Linden DEJ, Muckli L, Roth R, Drüen K, Ille N, Zanella FE, Singer W, Goebel R (1999) Combining MEG with fMRI in studies of the human visual system. In: Yoshimoto T, Kotani M, Karibe H, Nakasato N (eds) *Recent advances in biomagnetism*. Tohoku University Press, Sendai
- Schild HH (1990) *MRI made easy*. Schering AG, Berlin
- Scholz J, Klein MC, Behrens TEJ, Johansen-Berg H (2009) Training induces changes in white-matter architecture. *Nat Neurosci* 12:1370–1371
- Searle SR, Casella G, McCulloch CE (1992) *Variance Components*. Wiley, New York.
- Seehaus AK, Roebroeck A, Chiry O, Kim DS, Ronen I, Bratzke H, Goebel R, Galuske RAW (2013) Histological validation of DW-MRI tractography in human postmortem tissue. *Cereb Cortex* 23:442–450
- Setsompop K, Gagoski BA, Polimeni JR, Witzel T, Wedeen VJ, Wald LL (2012) Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planer imaging with reduced g-factor penalty. *Magn Reson Med* 67:1210–1224
- Sorger B, Peters J, van den Boomen C, Zilverstand A, Reithler J, Goebel R (2010) Real-time decoding of the locus of visuospatial attention using multi-voxel pattern classification, 16th edn. *Human Brain Mapping*, Barcelona
- Sorger B, Reithler J, Dahmen B, Goebel R (2012) A real-time fMRI-based spelling device immediately enabling robust motor-independent communication. *Curr Biol* 22:1333–1338
- Sorns O (2010) *Connectome*. Scholarpedia 5:5584
- Sorns O, Tononi G, Kotter R (2005) The human connectome: a structural description of the human brain. *PLoS Comput Biol* 1:e42
- Stejskal EO, Tanner JE (1965) Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *J Chem Phys* 42:288–292
- Subramanian L, Hindle JV, Johnston S, Roberts MV, Husain M, Goebel R, Linden D (2011) Real-time functional magnetic resonance imaging neurofeedback for treatment of Parkinson’s disease. *J Neurosci* 31:16309–16317
- Takeuchi H, Taki Y, Sassa Y, Hashizume H, Sekiguchi A, Fukushima A, Kawashima R (2010) White matter structures associated with creativity: evidence from diffusion tensor imaging. *Neuroimage* 51:11–18
- Talairach G, Tournoux P (1988) *Co-planar stereotaxic atlas of the human brain*. Thieme, New York
- Tournier JD, Calamante F, Gadian DG, Connelly A (2004) Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *Neuroimage* 23:1176–1185
- Tuch DS (2002) High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magn Reson Med* 48:577–582
- Tuch DS (2004) Q-ball imaging. *Magn Reson Med* 52:1358–1372
- van de Ven VG, Formisano E, Prvulovic D, Roeder CH, Linden DE (2004) Functional connectivity as revealed by spatial independent component analysis of fMRI measurements during rest. *Hum Brain Mapp* 22:165–178
- van de Ven VG, Formisano E, Roder CH, Prvulovic D, Bittner RA, Dietz MG, Hubl D, Dierks T, Federspiel A, Esposito F, Di Salle F, Jansma B, Goebel R, Linden DE (2005) The spatiotemporal pattern of auditory cortical responses during verbal hallucinations. *Neuroimage* 27:644–655
- Vapnik V (1995) *The nature of statistical learning theory*. Springer, New York
- Wager TD, Nichols TE (2003) Optimization of experimental design in fMRI: a general framework using a genetic algorithm. *Neuroimage* 18:293–309
- Wedeen VJ, Hagmann P, Tseng WY, Reese TG, Weisskoff RM (2005) Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magn Reson Med* 54:1377–1386
- Weiskopf N (2012) Real-time fMRI and its application to neurofeedback. *Neuroimage* 62:682–692
- Weiskopf N, Veit R, Erb M, Mathiak K, Grodd W, Goebel R, Birbaumer N (2003) Physiological self-regulation of regional brain activity

- using real-time functional magnetic resonance imaging (fMRI): methodology and exemplary data. *Neuroimage* 19:577–586
- Weiskopf N, Hutton C, Josephs O, Deichmann R (2006) Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: a whole-brain analysis at 3 T and 1.5 T. *Neuroimage* 33:493–504
- Worsley KJ, Marrett S, Neelin P, Evans AC (1992) A three-dimensional statistical analysis for CBF activation studies in human brain. *J Cereb Blood Flow Metab* 12:900–918
- Worsley KJ, Liao CH, Aston J, Petre V, Duncan GH, Morales F, Evans AC (2002) A general statistical analysis for fMRI data. *Neuroimage* 15:1–15
- Yacoub E, Harel N, Ugurbil K (2008) High-field fMRI unveils orientation columns in humans. *Proc Natl Acad Sci USA* 105:10607–10612
- Yeatman JD, Dougherty RF, Myall NJ, Wandell BA, Feldman HM (2012) Tract profiles of white matter properties: automating fiber-tract quantification. *PLoS One* 7:e49790
- Zimmermann J, Goebel R, De Martino F, van de Moortele PF, Feinberg D, Adriany G, Chaimov D, Shmuel A, Ugurbil K, Yacoub E (2011) Mapping the organization of axis of motion selective features in human area MT using high-field fMRI. *PLoS One* 6:e28716