Hence, on the whole, BOLD fMRI images are best understood as intermediary inscriptions whose function is, first and foremost, to bridge the otherwise insurmountable gap between the subject's active brain and the functional maps. As the output of the measurement procedure, BOLD fMRI images have a fixed material form. Owing to this fixed material form, they can be archived, copied and transported, shared within the scientific community and even reused in later studies.[269] However, as the following sections will make evident, the key feature of fMRI images is their mutability, which arises from the fact that various mathematical operations can be performed on them. Owing to their mutability, these images are able to fulfil their primary epistemic function as the working material for subsequent transformations. In what follows, we will examine these transformations and discuss their epistemic implications.

3.3 Preprocessing: Constituting the Analysability of fMRI Data

Having collected the imaging data for all their study participants, researchers then move on to the subsequent stages of the experiment, during which they process the raw datasets. Across these stages, researchers aim to translate the illegible and noisy fMRI datasets into visually accessible functional brain maps. Called the processing pipeline, this procedure entails a sequence of algorithmic steps that systematically address various types of noise. In the following sections, I will examine these steps by focusing on how researchers make judgments about what counts as noise in their data and which operations they perform to remove it. I will show that by making these judgments, researchers inscribe a range of both explicit and implicit theoretical assumptions into the imaging data. It is important to unpack these assumptions since they are invisible in the functional maps as the products of the analytical pipeline. Yet, although invisible, these assumptions inform the maps' potential scientific validity and their ability to produce new insights into hysteria or, at a more general level, any other phenomenon under study.[270]

Generally speaking, a processing pipeline comprises two distinct stages. Each stage is tailored to deal with a specific type of noise—random or systematic. The primary sources of random noise in an fMRI experiment include, first, brain processes unrelated to the experimental task, and second, variations in how the subjects performed the task at hand.[271] This type of noise is study-specific because it depends on the concrete experimental task and the subjects selected. To remove it, researchers deploy statistical analysis during the main stage of processing. But before statistical analysis can be

269 As discussed previously, the underlying structure of each slice is a matrix—an array of numbers arranged in rows and columns.

270 To demonstrate the analytical variability of fMRI processing pipelines, one meta-study focused on ten standard preprocessing and modelling steps. By considering between two and four default options for each step and then taking into account their various combinations, the authors arrived at 6,912 different pipelines. When applied to the same dataset, each pipeline resulted in a different functional map. See Carp, "Analytic Flexibility."

271 How a task is performed varies not just among different subjects but also over a single subject's repeated trials during the experiment. See Huettel, Song, and McCarthy, *Imaging*, 262.

used to translate them into functional maps, raw imaging data must first be prepared for analysis through preprocessing. The purpose of preprocessing is to remove non-meaningful changes in the MR signal caused by more or less predictable measurement constraints.[272] This type of noise is called systematic as it affects all fMRI studies independently of the task chosen.

Since systematic noise is not study-specific, its removal entails applying standard preprocessing steps. Therefore, many researchers tend to consider preprocessing less challenging than statistical analysis, which has to be tailored to each study.[273] As a result, researchers often report the preprocessing steps they implemented only summarily. For example, de Lange, Roelofs, and Toni described their entire preprocessing in a single sentence: "First, functional images were realigned, slice-time corrected, normalized to a common stereotactic space (MNI: Montreal Neurological Institute, Canada) and smoothed with a 10 mm FWHM Gaussian kernel."[274] However, in what follows, my analysis will show that researchers make far-reaching epistemic decisions at each of the steps listed above. More precisely, I will argue that preprocessing disciplines and standardises raw fMRI data by altering them to fit researchers'—often tacit—assumptions about what constitutes valid datasets for statistical analysis.

To perform preprocessing and the subsequent statistical analysis, researchers rely on specialised computer programmes. To begin with, they can choose among different software packages, most of which are freely available for research purposes. SPM, FSL, and AFNI are the most widely used open-source packages.[275] Significantly, although a shared analytical approach informs them, the programmes differ considerably in the sequence of the single steps, underlying theoretical concepts and mathematical modelling.[276] Besides, all packages are regularly updated with "substantial theoretical, algorithmic, structural and interface enhancements over previous versions."[277] Thus, both the differences across single packages and among various versions of the same software affect the outcome of processing.[278] Researchers are, therefore, obliged to specify which particular version of which software they used in their study. My analysis in the following will focus on the SPM—Statistical Parametric Mapping—which was the

272 Huettel, Song, and McCarthy, 267. By referring to imaging data as raw, I am merely emphasising that they are a direct output of the measurement and have yet to undergo preprocessing and statistical analysis.

273 Ashby, *Statistical Analysis*, 80.

274 De Lange, Roelofs, and Toni, "Self-Monitoring," 2053.

275 Developed by Karl Friston and colleagues, SPM is maintained by the Wellcome Trust Centre for Neuroimaging, University College London. FSL was created at the University of Oxford and AFNI at the National Institute of Mental Health in Maryland. See Poldrack, Mumford, and Nichols, *Handbook*, 8–9.

276 The different software packages predominantly rely on the general linear model (GLM) approach, which I will analyse in sections 3.4.1 and 3.4.2. For details, see also Poldrack, Mumford, and Nichols, 7–10.

277 http://www.fil.ion.ucl.ac.uk/spm/software/. Accessed on January 10, 2020. In the words of one of the SPM's developers: "The term 'SPM' does not really refer to a single piece of software, as many changes are made between each release." Ashburner, "SPM: A History," 792.

278 Carp, "Analytic Flexibility," 2, article 149.

first widely used software for fMRI analysis and continues to be the most popular.[279] Moreover, the SPM was used in our case study.

Notably, the analytical flexibility with which researchers can approach their data only begins with choosing the software. Each software version can be applied to the same dataset in immensely variable ways, both during preprocessing and even more so during statistical analysis. We will see that at each processing step, researchers can either choose among several pre-given standard options or define custom-made parameters. In doing so, they gradually construct a distinct chain of transformations tailored to the purposes of their study. Since these decisions have epistemic implications for the outcome of the processing, my analysis will examine how human judgment both guides and intervenes in the software-based operations throughout the chain of transformations that starts with raw imaging data and ends with functional maps. I will argue that the imaging data's mathematical and visual aspects fulfil distinctly different functional roles during this process. But before turning to the discussion of statistical analysis, in the following sections, I will first focus on illuminating the epistemic implications of the four major preprocessing steps: visual quality control, head motion correction, acquisition time correction and normalisation.[280]

In the remainder of this chapter, my analysis is based on close reading of fMRI studies of hysteria and multiple, more general publications that deal with the methodological aspects of functional neuroimaging. Importantly, my analysis is also substantially informed by practice-based insights I have gained while learning to use the SPM for fMRI data analysis. For this purpose, I participated in two courses for graduate students held by Dr. rer. nat. Torsten Wüstenberg at the Department of Psychiatry and Psychotherapy, Charité Campus Mitte Berlin in March 2014 and January 2015.

3.3.1 Identifying Visually Recognisable Noise

Strictly speaking, preprocessing comprises a sequence of algorithm-based steps.[281] Having selected the parameter settings at each preprocessing step, researchers let the software perform black-boxed mathematical operations on the fMRI slices. Since all transformations are conducted at the level of the numerical image matrix, it can be said that throughout preprocessing, fMRI images are treated as mathematical objects. This means that, at least in principle, researchers could clean their imaging data of systematic noise without even so much as glancing at them. However, standard textbooks on fMRI emphatically recommend that before submitting them to any algorithmic transformations, researchers should always look very closely at their

279 Poldrack, Mumford, and Nichols, *Handbook*, 8.

280 Preprocessing pipeline may comprise additional steps. For details, see Jenkinson and Chappell, *Neuroimaging Analysis*, 116–17, 122–30. I will not discuss such additional steps here, as they were not performed in our case study.

281 Recently, new methods have been developed that simultaneously combine all algorithmic preprocessing steps. Nevertheless, sequential preprocessing is still the dominant approach and will, therefore, remain the focus of my enquiry. See Jenkinson and Chappell, *Neuroimaging Analysis*, 121–22.

imaging data.[282] But what exactly can researchers see in the fMRI brain slices if these, as I have claimed, are visually illegible? Although it is impossible to determine the presence of task-induced brain activity by merely looking at fMRI brain slices, my analysis will demonstrate that, based on the visual appearance of the imaging data, researchers can make judgments about the data's tangential features. Specifically, we will discuss how researchers can assess the data quality by visually inspecting the fMRI slices.

Before they start processing them mathematically, researchers first examine the raw imaging data for potential measurement errors. This step is highly significant because, if overlooked, such measurement errors could lead to the creation of invalid functional maps. Typical sources of unwanted artefacts include potential technical problems with the scanner, various acquisition mistakes, errors in image reconstruction, and the experimental subjects' excessive head motion.[283] Researchers can use a range of automated software tools to check the quality of their data.[284] Yet, in addition to such quantitative examination, visual inspection of fMRI data on the computer screen is considered an indispensable part of quality control.[285] Many of the measurement artefacts listed above are visually discernible when the functional imaging data are viewed on the screen as grey-scale brain slices. Hence, it is considered that controlling the quality of fMRI images "is usually best done by eye, by just looking at the data."[286]

Indexically inscribed traces of various measurement artefacts can take different visual forms. For instance, some errors that arise from technical imperfections or scanner malfunction are visually detectable within single fMRI image slices. Such errors can appear as regularly repeating patterns of stripes or as unusual variations in the brightness spreading from the centre to the periphery of a 2D image.[287] Other artefacts take the form of a horizontal compression of the image towards the bottom or an unusual darkening of individual regions of a 2D slice. Less frequently, a shifted and warped version of the image may be superimposed on the original.[288] An experienced researcher can identify such visual distortions by merely glancing at a single fMRI slice. In other cases, the artefacts are not immediately apparent. Thus, to make the presence of an underlying anomaly visible, researchers must actively interact with the viewing software, for instance, by changing the default brightness setting.[289]

However, not all errors are detectable based on the inspection of single slices. More insidious artefacts are caused by unwanted changes that happen between the acquisitions of successive slices. Such errors become visually identifiable only when a time series of raw fMRI images are viewed in quick succession as a movie. To perform

282 Huettel, Song, and McCarthy, *Imaging*, 268; Jenkinson and Chappell, *Neuroimaging Analysis*, 89; and Poldrack, Mumford, and Nichols, *Handbook*, 37.

283 Huettel, Song, and McCarthy, *Imaging*, 267–68.

284 Huettel, Song, and McCarthy, 267–68.

285 Huettel, Song, and McCarthy, 267–68.

286 Jenkinson and Chappell, *Neuroimaging Analysis*, 89.

287 Huettel, Song, and McCarthy, *Imaging*, 268.

288 This particular artefact is called 'ghosting.' See Jenkinson and Chappell, *Neuroimaging Analysis*, 36, fig. 2.6.

289 Poldrack, Mumford, and Nichols, *Handbook*, 36.

such an inspection, researchers use various tools to animate all slices that constitute a single brain volume. In this way, they can examine the entire dataset, volume by volume, looking for rapid jerks in the animation or some other visual aspect that pops out of sequence.[290] Such visual disturbances are potentially significant, as they could point to a missing imaging slice or indicate that the experimental subject has abruptly moved the head during the measurement.

If they detect a visual anomaly in their data, researchers have to decide what further action to take. In some cases, they can remove the detected artefacts through mathematical processing and thus save the data. Yet, some measurement errors might be so extensive as to be beyond repair. In such cases, researchers have no choice but to exclude single slices, corrupt brain volumes or even an entire subject's dataset from further analysis.[291] Since the starting point of such far-reaching actions lies in the human inspection of the data's visual features, I argue that during preprocessing, various kinds of data visualisations are used operatively in the sense defined by Sybille Krämer. According to Krämer, when used operatively, visualisations function as tools that open new possibilities of actively engaging with and reasoning about the objects to which they refer. [292]

The above examples have shown that to look for potential traces of measurement errors in the data, researchers deploy different visual interventions, such as changing the brightness of individual slices or turning them into an animation. In doing so, they selectively articulate particular relations within the dataset and thus determine which kinds of artefacts are made visible in the form of particular visual patterns. Various artefacts might be present simultaneously in the same fMRI dataset. But to be visually brought forth and thus uncovered, each such artefact requires that the same dataset be visualised differently. It can, therefore, be said that various static and dynamic visualisations of the fMRI data are deployed during the quality control as flexible tools. Using these tools requires researchers to make active choices about how to visually configure their fMRI data to search for traces of possible acquisition errors, which would otherwise remain unnoticed. Significantly, such choices, in turn, enable researchers to classify the imaging data as either correct or corrupted.

Hence, although the fMRI data's numerical and visual forms contain the same information, they are not equivalent at the operative level. As we have seen, targeted visualisations can differentially display the pertinent relations in the data, which in the numerical form would remain inaccessible to researchers. Whereas the numerical form is crucial in enabling automated algorithms to transform the data mathematically, it is the visual form that addresses the human eye. In doing so, the data's visual form plays a central role in facilitating human judgments about the outcome of computer-based processes.

Although the process of visual quality control, as described above, may appear simple, it requires highly specific visual expertise. Functional imaging data are fuzzy and pixelated grey-scale images of brain slices. As I can testify from my experience, an

290 Huettel, Song, and McCarthy, *Imaging*, 268.

291 See, e.g., Espay et al., "Functional Tremor," 180.

292 Krämer, "Operative Bildlichkeit," 104–5.

untrained eye is unable to discern potential visual anomalies either in individual slices or in their animations. For this reason, researchers new to fMRI must first learn how to look for the visual features that could indicate underlying acquisition errors.[293] Novice researchers gradually acquire the visual expertise through practice by "repeatedly examining data from the same scanner."[294] The key aspect of this experiential learning is to develop implicit visual knowledge of "what 'good data' should look like."[295] In relation to what they know to be 'good data', experienced researchers can recognise pertinent visual distortions in a dataset. In other words, to differentiate between proper and corrupted data, researchers rely on an implicit comparison of what they have learned to see as salient visual features in a particular type of visualisation. Yet, although they can visually recognise such patterns and point to them on the computer screen, researchers are often unable to define them in verbally explicit terms.[296]

It appears to me that precisely the implicit character of researchers' expertise contributes to the ambivalent epistemic status of visual inspection in fMRI. On the one hand, the visual judgment of the human expert is accorded a crucial role in controlling and evaluating the output of the automated algorithmic processes. The relevant literature repeatedly advises researchers to visually examine not only the raw data following the acquisition but also the outcome of each preprocessing step to ensure that the algorithms did not accidentally introduce artefacts.[297] An expert human eye is thus deemed capable of identifying errors made by the 'blind' computer. But on the other hand, a visual inspection performed by a human expert is regarded as possibly biased and not entirely reliable unless complemented with automated calculations.[298] Moreover, by relying on their implicit expertise, researchers may recognise a visual indicator of an artefact. However, to pinpoint the exact source and the extent of the underlying problem and possibly remove it from the data, researchers must employ the software's algorithms. Whereas such algorithmic steps are typically reported in published articles, visual inspection remains unmentioned.[299]

***

Overall, this section has foregrounded the importance of visually examining the fMRI imaging data, especially during the initial quality control. I have emphasised how researchers' active and targeted engagement with different types of visualisations, both static and dynamic, and the researchers' implicit knowledge of what good data should look like underpin the process of visual data inspection. But I have also emphasised

293 For a pertinent analysis of how novice researchers acquire this kind of knowledge through embodied practice during training sessions with experienced colleagues, see Alac, *Digital Brains*, 67–145.

294 Huettel, Song, and McCarthy, *Imaging*, 268.

295 Huettel, Song, and McCarthy, 268.

296 Michael Polany has designated as 'tacit knowledge' the kind of knowledge "that cannot be put into word." Polanyi, *Tacit Dimension*, 4.

297 Huettel, Song, and McCarthy, *Imaging*, 272–73; and Poldrack, Mumford, and Nichols, *Handbook*, 35, 47.

298 Huettel, Song, and McCarthy, *Imaging*, 268; and Poldrack, Mumford, and Nichols, *Handbook*, 37.

299 See, e.g., Baek et al., "Impaired Awareness," 3; and Espay et al., "Functional Dystonia," 138.

that, despite its importance, visual inspection appears to be considered less 'objective' than clearly delineated algorithms. The reason for this, I suggest, is because the implicit knowledge that enables the visual judgment of the data's quality is neither quantifiable nor describable in clear-cut terms. It can only be transferred implicitly from researcher to researcher through the joint practice of working with and looking at images.

3.3.2 Erasing Temporal and Spatial Inconsistencies from fMRI Datasets

After passing the comprehensive quality control, raw fMRI data are submitted to two routine preprocessing steps—acquisition time correction and head motion correction. However, even deciding which of these two steps to perform first is a non-trivial matter. The problem is that, depending on the sequence of their application, these preprocessing steps could mutually interact, thus introducing errors into the data.[300] This fact alone already indicates that fMRI data undergo massive transformations during preprocessing. But what exactly happens to the images during these transformations, and what are the resulting epistemic implications?

Acquisition time correction targets temporal inconsistencies in the fMRI data caused by the sequential acquisition of 2D slices. For example, in the de Lange, Roelofs, and Toni study, each subject's brain volume was virtually divided into thirty-two slices collected sequentially over a period of 2.54 seconds.[301] This process was then repeated to acquire 547 brain volumes altogether. Due to this kind of acquisition, each slice in a single brain volume was collected at a different time point.[302] As a result, BOLD responses that occurred simultaneously across the brain were sampled at different stages of their temporal developments, depending on their relative spatial locations.[303] Yet, the problem is that the ensuing relative temporal displacement across slices counts as noise from the perspective of statistical analysis. This is because the underlying premise of the analysis is that BOLD responses in all slices within a single brain volume were measured simultaneously and that each two adjacent brain volumes were acquired at equidistant temporal intervals.[304]

To circumvent this problem, researchers submit fMRI data to the procedure called temporal interpolation during the acquisition time correction. This mathematical transformation enables researchers to use the actually measured data from neighbouring voxels to estimate the value of the MR signal that would have been obtained at each voxel had all the voxels in a single brain volume been sampled at once.[305] Importantly, to enable this calculation, researchers must first specify

300 Poldrack, Mumford, and Nichols, *Handbook*, 48.

301 De Lange, Roelofs and Toni, "Self-Monitoring," 2053.

302 Consequently, the most pronounced temporal delay is between the first and the last slice acquired in each volume, which in our case study amounts to 2.46 seconds.

303 "The slices acquired later in the volume show an apparently earlier response because the hemodynamic response has already started by the time that they are acquired." Poldrack, Mumford, and Nichols, *Handbook*, 41.

304 Sladky et al., "Slice-Timing Effects," 588–94.

305 Different mathematical methods can be used for combining the values from neighbouring data points to calculate the estimated signal value in each voxel. See Huettel, Song, and McCarthy,

the exact temporal order of the slice acquisition and then choose a reference slice from their dataset. As their reference slice, researchers can select the slice acquired at the beginning, halfway through the volume or at any other time point of the measurement.[306] The automated algorithms then temporally align all slices comprising a single volume to match the timing of the reference slice. They do so by shifting the sampling points (i.e., the value of the signal intensity measured) in all other 2D images, either forwards or backwards in time.

Significantly, at the end of the acquisition time correction, the spatial characteristics of the functional slices remain unchanged. Yet, the signal intensity measured initially at each voxel is replaced by a newly calculated numerical value. Hence, through this preprocessing step, the raw dataset with its temporally mismatching sequentially acquired slices has been transformed into a corrected dataset. This new dataset comprises a collection of brain volumes containing slices with a matching timing. Such mathematical modelling thus allows researchers to satisfy the requirements of statistical analysis by constructing a temporally consistent functional dataset.

Either before or after acquisition time correction,[307] the functional dataset must undergo an additional preprocessing step called head motion correction. This step aims to minimise a particularly vexing problem of image acquisition—the experimental subjects' unintended head motion, which could render the data unusable if excessive.[308] Although the subject's head is often fixed with padding during the data acquisition, it is nevertheless impossible to entirely avoid small-scale movements caused by an array of normal physiological reactions.[309] For example, subjects may reposition their shoulders due to tiredness, briefly hold their breath, or unintentionally move their head while performing the experimental task.[310] Crucially, even a displacement smaller than a millimetre changes the brain's relative position within the scanner's coordinate system, thus causing a misalignment between successively sampled brain volumes.[311] In such a case, the voxels with the same set of coordinates across subsequently acquired volumes no longer refer to the same location in the physical space of the brain. This, in turn, means that the same neuroanatomical structures occupy different locations across successive 3D fMRI images.[312] The resulting spatial mismatch violates the assumption

Imaging, 271. The SPM, however, does not offer researchers the possibility of a choice since the method called Fourier phase shift interpolation is hard-coded into the software. See Ashburner et al., "SPM12 Manual," 21–22.

306 Ashburner et al., "SPM12 Manual," 22–23.

307 Poldrack, Mumford, and Nichols, *Handbook*, 48.

308 Poldrack, Mumford, and Nichols, 44.

309 Huettel, Song, and McCarthy, *Imaging*, 272.

310 Even minimal head movements that arise from breathing and heartbeat cause motion artefacts referred to as physiological noise. However, if researchers choose to remove this particular type of noise, they have to deploy an additional preprocessing step, which I will not analyse here. For details on removing physiological noise from fMRI data, see Poldrack, Mumford, and Nichols, *Handbook*, 49–50.

311 Huettel, Song, and McCarthy, *Imaging*, 271.

312 It should be noted that apart from resulting in a spatial mismatch across fMRI volumes, head motion also additionally causes significant changes in the MR signal intensities stemming from misaligned voxels. In some cases, due to head motion, a portion of the brain might "move out of

of statistical analysis that "the brain is always in the same position" in images collected at different time points.[313] If uncorrected, this misalignment leads to incorrect functional brain maps.

To be able to erase the spatial mismatch between successive brain volumes, researchers must first estimate the head motion that caused it. Achieving this is far from simple because the subject's head motion arises from an individual interplay of many behavioural and physiological factors. In effect, the exact details of the brain's displacement during the acquisition remain necessarily unknown to researchers. Nevertheless, by employing computer algorithms to mathematically analyse the spatial mismatch across the collected images, researchers can derive assumptions about the brain's most likely position at each time point of the measurement. To do this, researchers must first choose a single fMRI volume from their dataset as a reference.[314] The automated algorithms then computationally superimpose all images in the dataset to this common reference and calculate the amount of each volume's misalignment. The brain is thereby treated as a rigid body—an object whose size and shape remain constant over the time of the data acquisition.[315]

Based on this assumption, the brain's presumed motion during the acquisition is modelled mathematically as a combination of three movements along and three rotations around the respective axes of the Cartesian coordinate system.[316] To obtain an estimate of the brain's motion, the black-boxed algorithms iteratively test different combinations of these six basic types of motion. They search for the combination that best describes the spatial mismatch between the reference image and the rest of the data. The goodness of fit of the estimate is determined mathematically by a quantity called cost function that measures how the intensities across different 3D images relate

---

the imaging volume, with an irreversible loss of data from the affected regions." Huettel, Song, and McCarthy, 271. And even if this does not happen, there are other problems. For instance, movements of the brain along the z-axis might cause some slices to "miss the [RF] excitation pulse, whereas others will experience two (or more) excitation pulses in rapid succession," thus leading to changes in "the relative BOLD signals recorded from each" of these slices. Ibid., 273–74. Moreover, the spatial displacement of the brain's magnetic field within the scanner's magnetic field elicits mutual interactions between these fields, producing unwanted field inhomogeneities. Finally, as a result of head motion, the locations of the brain's voxels in relation to the spatial encoding gradients necessarily change. All these changes induce distortions of the MR signals. See Jenkinson and Chappell, *Neuroimaging Analysis*, 118, box 3.5. Importantly, none of the motion-induced distortions of the MR signals can be removed through the deployment of head motion correction. Instead, additional processing steps have been developed that explicitly address this specific problem. But more often, and this is a point to which we will return later, motion-induced signal changes are filtered out during the stage of statistical analysis. For details, see Jenkinson and Chappell, 203–5.

313 Huettel, Song, and McCarthy, *Imaging*, 276.

314 Typically, the reference volume is a set of image slices acquired either in the middle or at the beginning of the measurement. Poldrack, Mumford, and Nichols, *Handbook*, 45. Alternatively, some studies compute the mean of the time series as the reference. See, e.g., Baek et al., "Impaired Awareness," 1626.

315 Poldrack, Mumford, and Nichols, *Handbook*, 45.

316 Poldrack, Mumford, and Nichols, 45.

to one other. Researchers can choose among different cost functions, each of which relies on a different mathematical model.[317]

Upon finished calculations, the algorithms construct a mathematical representation of how the subject's brain had presumably moved during the experiment. This mathematical representation is visualised by two sets of curves, which plot the brain's estimated displacements along and rotations around the respective Cartesian axes as the function of time (fig. 3.8).[318] Next, researchers can use the thus estimated motion to correct the spatial misalignment in the data. Having selected one of several available methods of spatial interpolation,[319] researchers use algorithms to calculate the data values that would have been acquired had the experimental subject remained motionless during the scanning.

First, the images are realigned (i.e., spatially transformed), which means that the original coordinates of the voxels are replaced by newly calculated ones. As a result of this operation, the 3D images are shifted from their native space (as determined by the measurement) into a newly defined image space.[320] After that, every 3D image is resliced—i.e., based on the values measured in the neighbouring voxels, the algorithms compute the signal intensities that would have been obtained at each new spatial point of the registered image.[321] In specialist terms, reslicing is referred to as 'bringing' or 'writing' the original image into the new image space.[322] Thus, in a two-step procedure, voxels are first shifted in place and then assigned new numerical values that designate the estimated signal intensities at the new locations.

As my analysis has shown, motion correction entails massive mathematical interventions into the spatial structure of the fMRI data. The native image space—i.e., the set of coordinates attributed to the imaging data by the measurement—is transformed into a 'corrected' image space, which is defined by newly calculated coordinates. The output of motion correction is a spatially more consistent dataset in which all fMRI volumes have been transformed to match the location of the reference volume. To ensure that this correction was performed accurately, researchers are recommended to inspect the dataset visually by viewing it as a movie.[323] If the correction has been successful, the resulting animation should be devoid of any jerky movements.

317 For details, see Jenkinson and Chappell, *Neuroimaging Analysis*, 169.

318 These estimations are stored additionally, as they play a role in statistical analysis. See Poldrack, Mumford, and Nichols, *Handbook*, 46. We will return to this point later in the chapter.

319 Different methods implement different mathematical relations between spatially neighbouring voxels to compute the estimated signal value. More accurate methods are computationally more demanding and thus take a considerably longer time to calculate. See Ashburner et al., "SPM12 Manual," 29. See also Poldrack, Mumford, and Nichols, *Handbook*, 46–47.

320 Jenkinson and Chappell, *Neuroimaging Analysis*, 160, 173–76.

321 Ashburner et al., "SPM 12 Manual," 29–32. See also see Jenkinson and Chappell, *Neuroimaging Analysis*, 176–79: and Poldrack, Mumford, and Nichols, *Handbook*, 28–30, 44–47.

322 Ashburner et al., "SPM 12 Manual," 29. See also Jenkinson and Chappell, *Neuroimaging Analysis*, 174–76. Jenkinson and Chappell use the term 'resampling' to refer to reslicing.

323 Poldrack, Mumford, and Nichols, *Handbook*, 47.

Figure 3.8. Visualisation of an experimental subject's estimated head motion during the fMRI data acquisition. From Ashburner et al., "SPM12 Manual," 259, fig. 32.5. ©Wellcome Centre for Human Neuroimaging, London.

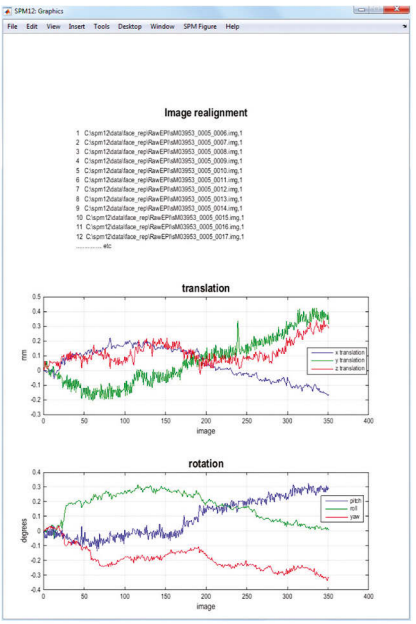

***

In sum, the combined aim of the preprocessing steps analysed above is to replace the signal intensities measured initially at respective time points and spatial locations with values that could not be sampled directly. We have seen that these interpolated values are necessarily estimates. Importantly, these estimates are not arbitrary. Instead, they are obtained by transforming the information contained in the original data through the application of standardised mathematical methods. As my analysis has underscored, all transformations are derived from a mathematical analysis of the original images. In effect, the algorithmic transformations recombine the initial signal measurements across the original images to generate the cleaned-up data. The algorithmic operations are black-boxed, with many of their aspects hard-coded into the software. Yet, I have shown that researchers make interpretational decisions throughout the process, such as choosing the reference image and selecting among the available parameter options, which include the type of cost function and interpolation method. These decisions are significant because each option entails different modelling strategies whose underlying

theoretical assumptions are inscribed into the new dataset. Thus, the adequacy of the steps chosen determines the potential accuracy of the outcome.

This extensive mathematical modelling serves to minimise the consequences of unavoidable technological and human-based measurement contingencies that introduce temporal and spatial inconsistencies into a single subject's fMRI dataset. In a group study, this procedure is performed separately for each subject's dataset. Through this procedure, each subject's newly calculated dataset is standardised and disciplined. In effect, it can be said that the implicit purpose of this standardisation is to mathematically approximate, as far as possible, an ideal situation, which no actual fMRI measurement can ever achieve. This ideal situation would entail generating a sequence of instantaneously acquired brain volumes from a motionless subject.[324] And although they cannot fulfil these ideal conditions, the corrected brain volumes—and the individual slices comprising them—are constructed as considerably more temporally and spatially consistent than those in the original raw dataset.

3.3.3 Establishing Anatomical Compatibility Across Data Types and Datasets

Once the temporal and spatial inconsistencies of each subject's functional dataset have been dealt with, the preprocessing moves to the subsequent stage. In this stage, researchers deploy two preprocessing steps specifically tailored to address multiple incompatibilities between different types of data and, in group studies, the inconsistencies across the individual subjects' datasets. In what follows, I will trace how the two designated preprocessing steps—coregistration and normalisation—standardise the imaging data. I will, in particular, foreground the epistemic implications of such standardisation.

In a single-subject study and many group studies, the next preprocessing step is coregistration.[325] Coregistration aims to enable the mapping of brain activations to anatomical locations after statistical processing has been completed. As mentioned previously, although fMRI images are not devoid of anatomical details, these are too imprecise to allow reliable identification of the brain's anatomical structures. This poses a significant problem since the aim of fMRI studies is to establish the anatomical locations of the task-induced brain activations of interest. To circumvent this problem, each fMRI study starts with acquiring a 3D high-resolution structural image that contains precise information about the subject's brain anatomy. However, although they refer to the same physical brain as the subject's fMRI images, structural slices are sampled with a different set of parameters. Hence, structural slices are characterised by a different spatial resolution, type of contrast, brain coverage, and

324 One caveat is that, as mentioned in footnote 312 above, motion correction cannot remove motion-induced changes in MR signal intensities from the fMRI data. Hence, even after this preprocessing step has been successfully applied, additional head motion artefacts remain in the data and must be dealt with during statistical analysis. See section 3.4.1.

325 In a single-subject fMRI study, coregistration is an individual step. As will become apparent shortly, in a group study, coregistration represents an optional substage of normalisation.

even artefacts.[326] Such differences make any direct comparison between structural and functional imaging data difficult, even when they stem from a single individual.

To combine the information contained in the two imaging modalities, researchers rely once again on computer algorithms. The use of algorithms enables the researchers to map the corresponding anatomical locations across functional and structural images through the process called coregistration.[327] Such algorithmic mapping is driven by a particular cost function chosen by researchers. The cost function quantifies the misalignment of the anatomical content between structural and functional images of the same subject by comparing pertinent image structures in both imaging modalities.[328] Through such quantitative image analysis, the algorithms estimate the parameters of the mathematical transformation that can best align the two different image spaces. By applying the transformation thus determined, fMRI images are realigned to match the image space of the structural image voxel-by-voxel. After that, researchers choose an interpolation method that uses the original data to compute the estimated signal intensities at the new locations.[329]

Through this chain of mathematical operations, coregistration constructs the spatial compatibility across the different imaging modalities. As a result, the anatomical information from the structural image can, at a later point, be deployed to anatomically designate the locations of activations in functional maps calculated from the fMRI images.[330] Using Ludwig Jäger's term,[331] we can thus say that during coregistration, researchers perform an intramedial transcription. They construct the anatomical legibility of the information obtained from fMRI brain slices by establishing a referential link to another type of image, i.e., structural imaging data.

Group studies, however, need to go beyond merely designating the anatomical locations of the experimentally detected activities in individual brains. Because they aim to produce generalisable results, group studies must combine data across multiple subjects. To enable comparison across subjects, researchers must first counter the problem that individual brains differ significantly. Notably, the individual differences

326 The differences in spatial resolutions and the types of contrasts that characterise these two imaging modalities were discussed in detail in sections 3.2.1 and 3.2.2.

327 Huettel, Song, and McCarthy, *Imaging*, 280–81. In fact, this step includes multiple operations since structural images have to be prepared for coregistration. Researchers first have to clean the images of various measurement artefacts, as well as algorithmically strip the brain of the skull and other non-brain tissue. They then proceed to segment the brain tissue into different types. These transformations rely on extensive mathematical modelling and require researchers to make interpretational decisions. For details, see Poldrack, Mumford, and Nichols, *Handbook*, 56–58.

328 The cost function typically used in coregistration is called boundary-based registration. It focuses on the boundaries between grey and white matter in both types of images while ignoring the rest of the visual content. See Jenkinson and Chappell, *Neuroimaging Analysis*, 212–13.

329 Ashburner et al., "SPM12 Manual," 43. See also Jenkinson and Chappell, *Neuroimaging Analysis*, 187–90.

330 This will be discussed in detail in section 3.5.1. At this point, it is important to emphasise that functional maps are devoid of any anatomical information and, therefore, cannot be coregistered directly onto structural images. For this reason, coregistration has to be performed with functional images. See also Jenkinson and Chappell, *Neuroimaging Analysis*, 170–71.

331 Jäger, "Epistemiology of Disruption," 72.

are not limited to the overall size and shape of each brain. Instead, they also include considerable variations in the positions and orientations of single anatomical structures across different subjects.[332] The crucial point is that brains of various shapes and sizes occupy arbitrarily different positions within the scanner's fixed coordinate system. Consequently, the same anatomical structures appear in divergent locations in images from different subjects and are thus designated by different sets of coordinates. Such inconsistencies hinder statistical analysis since automated algorithms can only calculate accurate group-level functional maps if the spatial coordinates of various neuroanatomical structures across all study participants are mutually aligned.[333]

To enable the comparison of fMRI datasets across individuals, researchers have to construct their mutual anatomical compatibility through a series of computerised steps jointly referred to as spatial normalisation. These steps transform each subject's image space—which is characterised by a contingent relation between that individual's neuroanatomical structures and the set of coordinates attributed to them through the measurement—into a shared space. In principle, spatial normalisation is similar to motion correction described in the previous section because it also mathematically transforms the imaging data to match them to a chosen reference image.[334] However, there are two crucial differences.

First, the underlying mathematical modelling in spatial normalisation is markedly more complex since the brain is no longer treated as a rigid body with a constant size and shape. During normalisation, the brain's size and gross anatomical structures are algorithmically transformed through "stretching, squeezing, and warping," thus substantially changing the geometry of the fMRI images in the process.[335] But although extensive, such spatial interventions are not arbitrary. Instead, they are limited by one crucial constraint—"an individual [anatomical] structure cannot be split up into separate structures and cannot disappear."[336] As in the previous processing steps, also in this case, researchers can select among various mathematical methods and levels of modelling complexity. Nevertheless, it is important to note that the software predetermines the range of available options of cost functions and interpolation methods researchers can choose.[337]

Second, unlike the preprocessing steps analysed so far, the reference image used in normalisation stems neither from the same fMRI dataset nor from the same measurement. When performing normalisation, researchers deploy an external reference image, which they can select from the software's various standard templates.[338] The most straightforward approach is to match the fMRI data to the software's standard functional template. Even though this approach is considered

332 Poldrack, Mumford, and Nichols, *Handbook*, 53.
333 Poldrack, Mumford, and Nichols, 17.
334 Huettel, Song, and McCarthy, *Imaging*, 282.
335 Huettel, Song, and McCarthy, 282.
336 Jenkinson and Chappell, *Neuroimaging Analysis*, 163.
337 Poldrack, Mumford, and Nichols, *Handbook*, 60–63.
338 Poldrack, Mumford, and Nichols, 59.

inaccurate,[339] many fMRI studies of hysteria—including the study by de Lange, Roelofs, and Toni—have implemented it. The more accurate but computationally considerably more elaborate approach is a multistep procedure. In the latter case, researchers first perform coregistration as described above and then align the subjects' structural images to one of the software's standard structural templates.[340] In both cases, the outcome of normalisation is a new fMRI dataset, whose image space matches the one defined by the template chosen.

All standard templates deployed by different software packages for fMRI processing are associated with one of the commonly used brain atlases. Their purpose is to provide what in the neuroimaging context is called a 'standard space.' That is, the templates offer a common 3D frame of reference in which a standardised set of Cartesian coordinates uniquely and consistently determines each neuroanatomical structure.[341] As opposed to the arbitrary positioning of the brain within the native space of each measurement, the standards space is defined by a fixed zero point and a fixed orientation of the coordinate axes in relation to particular anatomical landmarks.[342] For example, the zero point of the standard space is placed in the anatomical structure called the anterior commissure.[343] What happens during normalisation at the level of functional images is the following. The coordinates that the measurement had initially attributed to each voxel are translated into the standard space coordinates provided by the template. Ideally, through this translation, large anatomical structures across subjects should acquire the same set of standard coordinates by which these structures are uniquely determined in the given atlas.

In effect, the procedure of normalisation aims to homogenise the fMRI data by erasing the anatomical differences that characterise individual brains. In the process, all idiosyncratic anatomical features of an individual brain are treated as noise because they introduce spatial ambiguities into the data. Therefore, only by stripping each subject's dataset of individual anatomical specificities—and thus subsuming it to a standardised model—can the fMRI datasets of different subjects be made anatomically compatible. Such mathematically constructed anatomical compatibility is, in turn, a precondition for the mutual comparability of fMRI datasets across different subjects within a single study. Once they have been normalised, fMRI datasets of different subjects can be combined to compute group-level activation maps. Yet, at a more general level, normalisation of fMRI data also makes possible a direct comparison of imaging results across different studies. Specifically, "if data from two different studies have been normalized in the same fashion, then the areas of activity found in each study can be compared."[344] Hence, nowadays, even single-subject studies typically entail the step of spatial normalisation, as it facilitates the comparison of their results with other

339 "[A]lthough the overall outline of the brain will be accurate, structures within the brain may not be accurately aligned." Poldrack, Mumford, and Nichols, 59.

340 Poldrack, Mumford, and Nichols, 59–60.

341 Poldrack, Mumford, and Nichols, 54.

342 Poldrack, Mumford, and Nichols, 54.

343 Poldrack, Mumford, and Nichols, 54.

344 See Huettel, Song McCarthy, *Imaging*, 283.

studies.[345] Using Jäger's term,[346] it can be said that the anatomical consistency and the resulting mutual comparability of normalised fMRI datasets are constructed through their intramedial transcriptive transformation in relation to the software's standardised image templates.

There are two caveats, however. First, despite extensive mathematical modelling, the normalised fMRI datasets still retain residual anatomical differences. Hence, an additional preprocessing step called spatial filtering is often applied, which further reduces the residual anatomical differences by blurring the images.[347] Second, the concept of the standard space is not as stable or homogenous as it may appear at a superficial glance. Earlier neuroimaging studies deployed the Tailarach & Tournoux standard space derived from the identically named atlas.[348] This atlas is based on the dissection of a single hemisphere of a 60-year-old French woman's brain. However, the use of the Tailarach & Tournoux standard space is no longer considered "a good choice" in the neuroimaging community, as it is deemed unrepresentative of the general population and thus "provides a false sense of precisions and accuracy."[349]

For this reason, more recent studies have mostly relied on the template called MNI152 that was developed by the Montreal Neurological Institute (MNI) "as an average of structural MRI images from 152 young healthy adult subjects."[350] But the MNI152 is only the latest in several generations of MNI population-based templates, none of which are identical.[351] Moreover, because the MNI152 template is based on the brains of young, healthy subjects, it is unrepresentative of neurological patients.[352] Overall, my succinct overview has foregrounded that the standard space is a convention that continues to evolve with the ongoing research. The apparent consequence is that fMRI studies have implemented different standard spaces to align their data in the last two decades. This unavoidably resulted in inconsistencies in how researchers attributed anatomical locations to the activation patterns registered in their functional data.[353]

Finally, since there are no automated tools for assessing the quality of coregistration and normalisation, researchers are emphatically advised to visually inspect the results of the black-boxed mathematical operations that massively transform their data.[354] One way of doing it is to inspect the thus obtained volumes as a movie. Additionally, researchers can use various digital viewing tools to superimpose a single fMRI slice over the template and then "flick" back and forth between them to check if they sufficiently

345 See Huettel, Song McCarthy, 283. For a pertinent example, see Roy et al., "Dysphonia," 186.

346 Jäger, "Transcriptivity Matters," 50.

347 For details on spatial smoothing, as well as additional reasons why this preprocessing step is performed, see Poldrack, Mumford, and Nichols, *Handbook*, 50–52.

348 Poldrack, Mumford, and Nichols, 178.

349 For details, see Poldrack, Mumford, and Nichols, 177–78.

350 Jenkinson and Chappell, *Neuroimaging Analysis*, 191.

351 The initial MNI template was the so-called MNI305, with a lower resolution than the MNI152. See Poldrack, Mumford, and Nichols, *Handbook*, 55–56.

352 See Huettel, Song, and McCarthy, *Imaging*, 284.

353 Jenkinson and Chappell, *Neuroimaging Analysis*, 191.

354 Huettel, Song, and McCarthy, *Imaging*, 283; Jenkinson and Chappell, *Neuroimaging Analysis*, 183–84; and Poldrack, Mumford, and Nichols, *Handbook*, 65.

overlap.[355] Alternatively, they can extract the tissue boundaries from the template and overlay them on the normalised image to see how well they fit.[356] As in the case of visual inspection of raw imaging data, researchers have to learn through practice how to recognise potential artefacts and inconsistencies in their normalised imaging data.

***

My analysis in the last three sections has shown that, although considered to be the same for all experiments, preprocessing steps require researchers to make interpretational decisions about what counts as systematic noise in their datasets and which of the available transformation options to use to delete this noise. Automated algorithms then perform the chosen transformations at the numerical level of the imaging data. Yet, throughout my analysis, I have emphasised that the visual character of fMRI data nevertheless plays a crucial role during preprocessing. By interacting with the fMRI data's visual features, researchers determine if the automated algorithmic operations were carried out adequately. Moreover, we have seen that all these operations aim to reduce various idiosyncratic aspects of the measurement that introduced ambiguity into the data. Through these operations, fMRI datasets are mathematically constructed as increasingly mutually compatible.

Drawing on Latour, I argue that each preprocessing step is characterised by a trade-off between gain and loss.[357] What is lost at each step is the unwanted idiosyncrasy of the measurement, which arose either from the fMRI's technological limitations or from the experimental subjects' behavioural and physiological contingencies. My analysis has underscored that this deletion is performed under specific constraints. The images are transformed first by shifting the voxels to locations defined by new sets of coordinates. Then the corresponding signal intensities at these new locations are calculated by using the values from the neighbouring voxels. The values thus computed are only estimates of the data that would have been collected in an unattainable situation, which would have allowed the instantaneous acquisition of successive fMRI volumes from a static brain of a standard size and shape. Nevertheless—and this is crucial—the use of the Cartesian coordinate system and a particular set of mathematical operations ensure that the transformation of the original raw dataset into a corrected one is traceable, at least in principle.[358] Provided that they did not result in errors, the mathematical operations retain an unbroken referential link to the original signal,[359] which, in turn, is indexically related to the individual subject's active brain.

Conversely, what is gained through preprocessing is the temporal, spatial, and anatomical consistency within and across the newly calculated datasets. Through

355 Jenkinson and Chappell, *Neuroimaging Analysis*, 183–84.

356 Jenkinson and Chappell, 183–84.

357 Latour, *Pandora's Hope*, 70–71.

358 It should be noted that all interpolations "involve some degradation of the image, as some information from the original image is lost." Jenkinson and Chappell, *Neuroimaging Analysis*, 178. Put simply, the price researchers pay for deleting systematic noise is a partial loss of potentially meaningful information.

359 I am using the term 'referential' in Latour's sense. See Latour, *Pandora's Hope*, 71–72.

algorithmic operations of mutually aligning the fMRI images to one another, as well as matching them to other imaging modalities and external image-based templates, researchers create a dataset that is "compatible with already-established centres of calculation."[360] Importantly, the output of these transformations are 4D functional datasets that are still illegible—when preprocessed fMRI datasets are submitted to visual inspection, even experts cannot 'read' them. In short, by looking at these images, it is still impossible to determine which voxels exhibit task-induced activity and which do not. Nevertheless, thus standardised, the images can now finally undergo statistical analysis that will translate them into legible brain maps. Hence, as shown by my analysis, the purpose of preprocessing is to construct the analysability of the fMRI datasets while at the same time preserving their indexicality via a chain of traceable mathematical operations.

3.4 Statistical Analysis: Articulating the Task-Induced Neural Activity of Interest

Preprocessed functional 4D datasets remain illegible because the pertinent information concerning the brain activity of interests they entail is still spread across multiple brain volumes and buried under random noise. To construct the legibility of their fMRI data, researchers must determine which areas of the subjects' brains can be declared active. They do this by using statistical analysis, which enables them to make judgments about the "underlying patterns in the data" ridden with random noise.[361] Instead of more commonly known descriptive statistics that merely summarise the data, fMRI studies apply inferential statistics. This type of statistics permits researchers to use the datasets from their subject sample to make claims about a larger population.[362]

Inferential data analysis is based on the process called hypothesis testing. Generally speaking, this type of statistical analysis starts with the formulation of two opposing claims—the null hypothesis and the alternative hypothesis.[363] In the subsequent step, statistical tests are used to evaluate which of the two hypotheses describes the data with a higher probability. In fMRI, the null hypothesis amounts to the claim that the task had no effect on the data, or in other words, that there is no temporal correlation between the variation in the BOLD time series and the different experimental conditions. The alternative hypothesis states that the measured differences in the BOLD signal's average intensities between the task and the control condition are temporally correlated with the experimental intervention.[364]

During hypothesis testing, the analysis software executes automated statistical tests for each voxel independently. This voxel-by-voxel approach is known as mass

360 Latour, 71–72.
361 Worsley, "Statistical Analysis," 251.
362 Worsley, 251.
363 Huettel, Song, and McCarthy, *Imaging*, 331.
364 Huettel, Song, and McCarthy, 331.