
In summary, the two consecutive studies by de Lange, Toni, and Roelofs generated categorically different imaging findings through the applications of two different analytical approaches to the same fMRI dataset. We have seen that each of the two approaches was informed by a substantially different model of the brain function. In one case, the focus was on strictly localised activations (functional segregation), whereas in the other, on the dynamic connections among spatially remote brain areas (functional integration). Just as significantly, each approach also rested on partly contrary definitions of what counted as the information of interest in the fMRI data instead of noise. Therefore, each approach required that the researchers deploy different kinds of mathematical transformations to obtain what they defined as pertinent information.

In effect, my analysis has shown that the kind of information that is articulated from a particular fMRI dataset and translated into a legible statistical map is, at the most basic level, predicated on the model of the brain's functional organisation which underpins the analytical approach chosen by researchers. Because these models are not mutually exclusive, they can be applied in separate analytical procedures to the same fMRI dataset to construct multiple, mutually complementary statistical brain maps. Through the use of such mutually complementary analyses, a single fMRI dataset is constructed as what I would like to designate as *semantically multipotent*. What I mean by this is that each fMRI dataset holds the potential to be made legible in multiple epistemically valid ways. As we have seen, it is up to researchers to decide which specific semantic potential of their fMRI dataset they want to articulate to answer their study-specific research questions. In each case, the result of such an articulation is a particular statistical brain map.

3.5 Visualising Functional Brain Maps: Ascribing the Symbolic Meaning

Only after they have completed all the steps entailed in the time-consuming data analysis and thus obtained the statistical maps of their choice can researchers finally turn to evaluating the empirical results of their experiment. To put it more plainly, it is not before this point that researchers can even see which brain areas were differentially activated—with sufficient statistical significance—by the comparisons of the experimental conditions they chose to test. Having invested weeks or even months into painstakingly constructing their functional maps, researchers can, at last, use them to answer two crucial questions. In which anatomical regions of the brain did the experimental intervention trigger neural responses? And, how do such patterns of brain activity relate to cognitive processes that play a role in the formation and manifestation of the hysterical symptom of interest, or more generally, any other phenomenon under investigation?

Answering these questions requires researchers to make sense of their statistical brain maps. Yet, there is one crucial point that I want to make. Although the statistical brain maps are legible, their exact informational content and medical meaning are far

from obvious even to an expert.⁴⁸² As I will argue in the remainder of this chapter, the meaning of the maps has to be constructed in a step-by-step procedure. We will see that during this procedure, different visualisations play productive roles in allowing researchers not only to understand their maps but also, in the final instance, to arrive at a particular interpretation. In short, of central concern to our discussion is what kinds of visualisations researchers use during this procedure and how they interact with these visualisations.

In what follows, I will first analyse how researchers deploy highly interactive digital visualisations to examine the maps and make them interpretable in anatomical terms. In the subsequent section, I will return to the case study at the centre of this discussion to examine how researchers visually fix their results in the form of publishable composite figures that—as Martina Merz fittingly formulated it—“travel well” within the research field.⁴⁸³ Finally, drawing further on the case studies, I will show how by constructing a complex network of intermedial and intramedial references,⁴⁸⁴ researchers institute their fMRI figures into symbolic signs of cognitive phenomena. I will argue that in doing so, researchers are able to develop hypotheses about the potential neurocognitive basis of hysterical symptoms that they study.

3.5.1 Utilising Visualisations to Explore and Assess the Empirical Results

As discussed previously, a statistical brain map that researchers have created through hypothesis testing of a chosen contrast of experimental conditions and then corrected for multiple comparisons is, in effect, a 3D collection of active voxels. Moreover, we have seen that only those voxels—or clusters of voxels—were declared active whose calculated levels of statistical significance survived the corrected threshold. Hence, in the resulting statistical brain map, each active voxel contains a numerical value determined by the test statistic calculated for the chosen contrast at a given location. Conversely, inactive voxels are empty because, after thresholding and the multiple comparisons correction, their numerical value has been set to zero. Thus, a statistical map is, in essence, a collection of spatially organised quantitative information. Yet, as soon as the calculations underlying the map's creation are finished, the software automatically transforms the resulting quantitative information into multiple visualisations. In what follows, on the example of the SPM software, I will analyse how researchers work with such visualisations to assess the quantitative results of their experiment by making judgments about the anatomical locations of the brain activities identified. I will show that different ways in which the fMRI maps are visualised during this working process play crucial roles in facilitating the researchers' ability to 'read' these maps with sufficient accuracy.⁴⁸⁵

482 By designating the maps as 'legible,' I am foregrounding that the information of interest (i.e., the location of activated voxels) has become accessible to visual inspection.

483 Merz, "Designed for Travel," 349–50.

484 Jäger, "Transcriptivity Matters," 50.

485 I am using the term 'reading' here in Krämer's sense to denote the learned ability to overlook the epistemically insignificant visual features while also knowing which relevant visual features to focus on to obtain the information of interest, which is encoded in the image. See Krämer, "Operative Bildlichkeit," 102.

The first visualisation that the SPM automatically generates upon the completed hypothesis testing and thresholding for a given contrast is what, in specialist terms, is called the maximum intensity projection (fig. 3.10, top).⁴⁸⁶ This composite visualisation is also fittingly referred to as the glass brain views. As suggested by the latter designation, in this visualisation, the brain is treated as a transparent object and shown simultaneously in three mutually orthogonal planes of the Cartesian coordinate system. Each of the three views displays a grid pattern, which is overlaid with an outline of the brain in the sagittal (longitudinal), coronal (frontal), and axial (horizontal) planes, respectively. Statistically significant activations are grey-scale coded and projected through the brain along the given viewing axis onto each outline. Notably, only the peak activation along each viewing axis (i.e., a single voxel with the highest numerical value) is made visible in each respective plane.⁴⁸⁷ Conversely, all other statistically less significant activations along the given projection axis remain invisible. Hence, each of the three mutually orthogonal planes provides only a partial picture of the 3D activations. The red arrow, which appears in each 2D view, points to the same spatial location across the three planar projections. By left-clicking and then dragging this arrow, researchers can move it to a different location within one of the glass brain views and thus actively explore the spatial distribution of the activations projected. Since the visualisation is interactive, moving the arrow in one view leads to the automatic readjustment of the respective positions of the corresponding arrows in the other two outlines.

Importantly, the glass brain views simultaneously display peak activations located not just on the surface but also in the deeper structures of the brain. Thus, the major advantage of this composite visualisation is that it enables researchers “to see all of the [peak] activations at once.”⁴⁸⁸ In other words, the glass brain views provide researchers with a global visual overview of the results. However, the glass brain views have one major caveat—working with them is far from simple. Since these empty brain outlines lack anatomical landmarks, researchers require considerable expertise to be able to judge the location of activation of interest. What is even more challenging is that the individual glass brain views are undetermined if viewed in isolation. Put simply, many different 3D spatial distributions of the activations could result in exactly the same 2D projections along the axes.⁴⁸⁹ To even approximately localise the activations, researchers must learn to mentally combine all three outlines by relying on the red arrows as the points of orientation across the views. In short, by integrating the partial information displayed in the separate 2D views, researchers have to build up a mental picture of 3D activations. Acquiring such a visual skill requires extensive practice.


486 Ashburner et al., “SPM12 Manual,” 248. See also Poldrack, Mumford, and Nichols, *Handbook*, 175–76.

487 This explains why this type of visualisation is called the maximum intensity projection. For details, see Wallis and Miller, “Three-Dimensional Display,” 535–36; and Wallis et al., “Three-Dimensional Display,” 297–98.

488 Huettel, Song, and McCarthy, *Imaging*, 371.

489 Huettel, Song, and McCarthy, 371.

Figure 3.10. Top: glass brain views of statistically significant activations computed for the contrast of the experimental conditions designated in the design matrix to the right. Bottom: statistical table listing all clusters of activations above the chosen level of significance. From: Ashburner et al., “SPM12 Manual,” 291, fig. 33.5. ©Wellcome Centre for Human Neuroimaging, London.



Despite this caveat, experienced researchers, who know how to skilfully read the glass brain views, can deploy these visualisations as highly effective tools for the initial assessment of the experimental results. They can use these images to judge how much activation was induced by the given contrast across the brain. Moreover, skilled researchers can roughly bring different activations into spatial relations to one another by navigating the glass brain with the help of the red arrows. Conversely, if a contrast of interest resulted in blank glass brain views, researchers can choose among several possible courses of action. For instance, they can conclude that the lack of activation in the given map is meaningful. This is precisely what de Lange, Roelofs, and Toni did for some of their contrasts.⁴⁹⁰ In such cases, the absence of statistically significant differential activations is taken to indicate that the contrasted task conditions induced the same neural effects. Alternatively, researchers can also decide that their empty or

490 De Lange, Roelofs, and Toni, “Self-Monitoring,” 2054.

almost empty glass brain views mean that the corrected thresholds they used were too stringent. At this point, they might choose to recalculate their maps in order to lower the threshold or to revert to working with uncorrected maps.⁴⁹¹ Although not uncommon in practice, these two latter approaches are intensely criticised in the neuroimaging literature. The general consensus is that both of these approaches increase the false positive rates and thus lead to erroneous interpretations of empirical results.⁴⁹²

In addition to the glass brain views, the SPM simultaneously generates a supplementary visualisation. In this visualisation, the same set of results is displayed on the computer screen as a table containing numerical values (fig. 3.10, bottom). This table effectively summarises all relevant statistical information entailed in an fMRI brain map by organising them into rows and columns according to different categories. Among other information, individual columns contain the numerical values of the calculated test statistics, corrected and uncorrected significance values for each cluster of activation, and the set of 3D coordinates that determine the locations of the peak activations within each cluster listed.⁴⁹³ Just like the glass brain views, this table is also interactive. Hence, by clicking on a row of coordinates that denote a specific cluster of interest, researchers can inspect its various statistical values in more detail. Furthermore, the table and the glass brain views are mutually interlinked. Clicking on a set of coordinates in the table causes the red arrows in the glass brain views to move to the corresponding location.

This interlinking across visualisations is highly significant, as it aids researchers in aligning and mutually combining the glass brain views with the statistical table to gain a more comprehensive understanding of their statistical brain map. The table provides researchers with a summary of the map's underlying quantitative information, which they use to evaluate the statistical relevance of the activations identified. Yet, based on the table alone, it would be difficult to comprehend the spatial distribution of the activations, whose locations in this type of visualisation are denoted exclusively by sets of coordinates. Researchers, therefore, combine the statistical table with the glass brain views, which foreground the spatial relations among the activations at the expense of the quantitative information. I thus argue that the statistical table and the glass brain views are two types of visualisation that provide mutually complementary perspectives on the same statistical map. Each of them visually articulates a different aspect of the same map by foregrounding either its quantitative or spatial character. Since both of these aspects are crucial for making sense of the information contained in the map,

491 See, e.g., Stone et al., "Simulated Weakness," 963, 965.

492 Poldrack et al., "Scanning the Horizon," 121–22. In fact, such approaches are viewed as instances of p-hacking, a questionable practice of actively seeking and thus artificially inflating the statistical significance of the empirical results by manipulating the data. In addition to using uncorrected thresholds or recalculating the statistical maps, other instances of p-hacking include exploring "various analytic alternatives [during the stage of statistical analysis], to search for a combination that yields 'statistical significance,' and to then report only what 'worked.'" Simmons, Nelson, and Simonsohn, "False-Positive Psychology," 1359. On problems related to p-hacking, see also Head et al., "P-Hacking."

493 Ashburner et al., "SPM12 Manual," 250–51.

these two interactive visualisations effectively supplement each other. Jointly, these visualisation tools enable researchers to explore their empirical results.


Having gained a global impression of their map, researchers can then zoom in on single clusters of activation to inspect local neural responses that were elicited by the contrast of experimental conditions for which the map was computed. For this purpose, the SPM offers the possibility of visualising the mean estimated effect sizes in the form of a bar diagram (see fig. 3.14b). Each separate bar in the resulting diagram designates the estimated effect size for a particular experimental condition comprising the contrast. As discussed previously, the test statistic at a given voxel—i.e., the numerical value contained in the map—quantifies the statistical significance of the local BOLD response to that contrast. The estimated effect sizes provide supplementary information about the calculated strengths of the individual responses elicited by the experimental conditions that make up the contrast.⁴⁹⁴ Researchers can also visualise the fitted BOLD response at a single voxel to examine how the signal from that location changed throughout the measurement. The thus visualised curve displays the time course of the BOLD response predicted by the design matrix and then fitted to the data from that voxel during the stage of model estimation (fig. 3.11).⁴⁹⁵ Importantly, however, besides the fitted curves, the visualisation also displays the preprocessed time course of the signal that was actually measured at that voxel. Such simultaneous visualisation of the fitted curve and the actual data enables researchers to visually assess the quality of their GLM-based study-specific model.

In essence, both the bar diagram and the fitted BOLD response are derived from intermediary inscriptions that, as analysed previously, partook in the process of creating the statistical maps. By visually examining these supplementary visualisations, researchers can evaluate the quality of the steps through which the map was produced. The fact that researchers actively inspect these intermediary inscriptions while assessing the maps is significant. It demonstrates that the epistemic status of statistical brain maps is predicated on their continued dynamic embeddedness into the chains of transformations underpinning their production. Hence, to adequately evaluate the empirical results obtained through the brain map, researchers have to perform several interrelated operations. First, they have to combine spatial and numerical visualisations of the maps. Just as importantly, they also have to examine visualisations that provide both global and local overviews of the findings. Moreover, researchers need to be able to inspect previous inscriptions along the chain of transformations through which the resulting maps were constructed. Thus, I argue that all these highly versatile, mutually interlinked types of visualisations are required to enable researchers to examine the statistical maps from different perspectives. In effect, it is such a complementary use of multiple visualisations that makes the experimental findings in their complexity graspable to researchers.

494 As discussed previously, each contrast entails a comparison (i.e., a subtraction) of two or more experimental conditions. See Ashburner et al., “SPM12 Manual,” 251–52.

495 Ashburner et al., 251–52. See section 3.4.2 for a discussion of the model estimation.

Figure 3.11. The red line visualises the modelled BOLD response at a given voxel, whereas the grey line shows the preprocessed time course of the signal measured. The dots designate the individual sampling points.



After the initial assessment of their empirical results, researchers then proceed to identify the anatomical locations of the activated clusters of voxels. As discussed previously, neuroimaging research operates under the premise that distinct brain areas have specialised functions. This means that inferences about the underlying neural basis of hysterical symptoms can only be made in relation to concrete neuroanatomical structures. Yet, the problem is that, as shown earlier, statistical maps are devoid of any anatomical information. Instead, the relative spatial locations of the activated voxels are designated by respective sets of the standard space coordinates. Hence, to enable the anatomical localisation of the activations, the standard space coordinates must be brought in relation to brain anatomy. This is done by overlaying the statistical map onto another image that displays brain anatomy while paying particular attention that the coordinates of the statistical map and the anatomical image are mutually aligned. A variety of anatomical images can be used for this purpose. But as I will show in what follows, choosing which particular type of anatomical image to deploy is epistemically significant because each type differently configures the legibility of the superimposed statistical map. Specifically, we will see that the choice of anatomical images shapes not only how researchers work with the statistical map but also what they can see in it.

The most common approach to anatomically visualising the clusters of activation is to superimpose the statistical map onto 2D grey-scale anatomical images.⁴⁹⁶ To make it stand out against the grey-scale base image, the statistical map is colour-coded. In other words, different numerical values of the active voxels' test statistics are ascribed different colours. The SPM and comparable software packages offer various default colour-coding options, including the commonly used red-orange-yellow scale or the

496 At this point, researchers can continue to use the SPM. Alternatively, they can revert to other free programmes—such as MRICron or FSLView—which were specifically developed for visualisation purposes and are thus more flexible. This comment is based on my experience as a participant in SPM courses at the Department of Psychiatry and Psychotherapy, Charité Campus Mitte Berlin in March 2014 and January 2015.

rainbow colours. In general, darker colours denote lower, whereas the brighter refer to higher levels of statistical significance.⁴⁹⁷ But it is important to note that the concrete ascription of colours is entirely arbitrary. The colour-coding fulfils a purely utilitarian function as it allows researchers to distinguish different levels of statistical significance by merely glancing at the map.

Various types of anatomical images in different spatial orientations can serve as the base for displaying the activations. The SPM offers the possibility of using canonical anatomical templates of a standard brain in MNI space, which I mentioned earlier while discussing the normalisation. However, this option is considered inaccurate. The reason is that the standard brain cannot account for individual anatomical differences across subjects even after their brains have been normalised to this template.⁴⁹⁸ In single-case studies, the most accurate approach entails using the subject's own structural scans that were coregistered to the functional data during preprocessing. Conversely, a group-averaged map is ideally projected onto a mean structural image obtained by averaging the normalised anatomical scans across all group members.⁴⁹⁹ At a superficial glance, the mean anatomical image might appear imprecise because the averaging unavoidably results in the blurring of anatomical structures. Yet, somewhat paradoxically, this blurring "accurately reflects the imprecision in the functional data due to underlying anatomical variability."⁵⁰⁰ Thus, displaying the group activation on an anatomically more precise image, such as a standard template or an individual subject's structural scans, is considered to misrepresent the anatomical imprecision of the functional map and, in turn, lead to potential anatomical mislabelling of group activations.

Having decided which anatomical image to use as a base, in the next step, researchers can choose among different viewing options. They can either overlay the activations on three adjacent horizontal slices or, similarly to the glass brain views, on three mutually orthogonal sections along the respective axes of the Cartesian coordinate system.⁵⁰¹ In both cases, the identical location in all three simultaneously visible viewing planes is signified by an interactive crosshair—a point at which two orthogonal lines intersect. The key advantage of using the orthogonal sections is that they allow researchers to virtually 'move' through the entire brain volume along each axis (fig. 3.12). By selecting a different set of coordinates, researchers can shift their position within the virtual brain to another location. The new location is visualised on the screen by a new set of mutually orthogonal 2D sections. By repeating this operation, researchers can actively engage with the visualised map to explore the anatomical locations of different activation clusters.

This dynamic working process serves to circumvent the fact that the slices reveal only those activations that are located within the visualised planes, whereas all the rest of the activated clusters remain occluded. On their own, such partial views are insufficient because clusters of activation are 3D and thus spread across multiple

497 Huettel, Song, and McCarthy, *Imaging*, 369.

498 Devlin and Poldrack, "Tedious Anatomy," 1035.


499 Devlin and Poldrack, 1037.

500 Devlin and Poldrack, 1037.

501 Ashburner et al., "SPM12 Manual," 252–53.

anatomical structures. To gain the impression of each cluster's exact 3D shape and find out which anatomical regions it encompasses, researchers must navigate the virtual brain and visually inspect its multiple locations. Hence, these dynamic composite visualisations that fuse structural images with a functional map and allow a self-directed movement through the virtual brain are used as explorative tools. Researchers actively deploy these visualisations to make sense of their empirical results.

Figure 3.12. SPM's anatomical visualisation of an fMRI activation map in the form of 2D sections. From: Ashburner et al., "SPM12 Manual," 253, fig. 31.19. ©Wellcome Centre for Human Neuroimaging, London.



Additionally, researchers may choose to project the maps onto a 3D brain rendering that the software can compute from the structural MRI data.⁵⁰² Like 2D visualisations of the brain, a 3D surface rendering is also interactive and can be rotated on the computer screen and viewed from different directions. But unlike single 2D slices and sections, 3D renderings show only the activations located on the surface of the brain. Consequently, in such a visualisation, those active clusters that occupy internal cerebral structures necessarily remain hidden from view. Nevertheless, the significant advantage of this type of visualisation is that it provides 3D spatial views of the brain's anatomical structures. Such views are considered visually more graspable than 2D slices or sections (fig. 3.13).⁵⁰³ That is, even for an expert, it appears to be easier to visually

502 Ashburner et al., 253–54. See also Poldrack, Mumford, and Nichols, *Handbook*, 176–77.

503 Devlin and Poldrack, "Tedious Anatomy," 1037. See also Wandell, Chial and Backus, "Visualization and Measurement," 739.

differentiate among various cerebral structures and to identify the spatial distribution of statistically significant activations when viewing the brain displayed as a 3D object. Hence, the choice of a particular type of anatomical visualisation decisively influences the graspability of the anatomical locations of neural activations.

However, there are also disadvantages to using 3D surface renderings. First, the process of rendering a 3D structural image can be computationally very intensive and time-consuming. The second and far more serious problem arises from the characteristics of the brain anatomy. The cortical surface is a highly folded structure that comprises an undulating pattern of ridges (i.e., gyri) and grooves (i.e., sulci).⁵⁰⁴ Moreover, folding patterns are highly individual and thus vary considerably across different individuals.⁵⁰⁵ Due to such variations, 3D surface models rendered from group-averaged structural images are blurred and, therefore, anatomically imprecise. How to anatomically map the group-level activation patterns onto such 3D brain models with sufficient accuracy is far from straightforward and can differ considerably between software packages.⁵⁰⁶ Depending on how a particular software performed this operation, called surface-based registration, the same activation pattern can be attributed to distinctly different anatomical locations.⁵⁰⁷ Finally, the third problem with using 3D renderings is that “much of the cortical gray matter is buried within sulci.”⁵⁰⁸ Consequently, in such a visualisation, not just the activations that occupy internal structures but also all the activations located within deep sulci necessarily remain hidden from view and thus inaccessible to visual inspection (fig. 3.13, right).⁵⁰⁹

Irrespective of the specific advantages and disadvantages that a choice of a particular structural base image entails, the shared purpose of all such visualisations is to enable the anatomical localisation of statically significant activations. Experienced researchers may be able to accurately label anatomical structures through careful visual inspection of the statistical maps thus visualised.⁵¹⁰ Otherwise, researchers use automated software tools that perform the localisation by segmenting the underlying

504 Huettel, Song, and McCarthy, *Imaging*, 189.

505 For details, see Jenkinson and Chappell, *Neuroimaging Analysis*, 223–29.

506 I am grateful to Torsten Wüstenberg for pointing this out to me.

507 For a detailed description of the challenges involved in this operation, see Jenkinson and Chappell, *Neuroimaging Analysis*, 227–29.


508 Wandell, Chial, and Backus, “Visualization and Measurement,” 739.

509 To circumvent this particular problem, researchers may choose to display the activations either on so-called ‘inflated brains’ or on flattened cortical surfaces. Both ‘inflated brains’ and flattened surfaces are computed by mathematically deforming the initial 3D rendering of the brain “to allow for better visualization.” Jenkinson and Chappell, *Neuroimaging Analysis*, 100. The mathematical transformation entailed in obtaining an ‘inflated brain’ “acts in much the same way as taking a crumpled paper bag and blowing air into it: the bag will inflate and the overall surface will become smoother.” Ibid. As a result of this mathematical transformation, the activations that had thus far remained hidden within the sulci would become visible in the inflated brain. “To continue the analogy, you could then flatten the bag by making some cuts down its side and by pressing it flat on a table.” Ibid. The result of this second operation is a flattened cortical surface. For complex mathematical modelling required to compute such visualisations, see Wandell, Chial, and Backus, “Visualization and Measurement,” 741–51.

510 Devlin and Poldrack, “Tedious Anatomy,” 1036.

structural image into standard anatomical parcellation schemes. Called automated anatomical labelling, the latter approach is entirely black-boxed and not consistently “accurate across individuals” with highly variable brain anatomies or across different brain regions.⁵¹¹ Therefore, when using the automated approach, researchers are advised to verify the quality of the thus obtained results by visually comparing them to one of the anatomical brain atlases commonly used in neuroimaging.

Figure 3.13. Comparative views of 2D (left) and 3D (right) anatomical visualisations of the same fMRI map.



In fact, despite the increasing popularity of automatic labelling, the relevant literature still recommends that, either instead of or in addition to deploying the available automated tools, researchers should manually determine the anatomical location of the activation. Using a brain atlas as a reference, they should rely on visual comparison to identify pertinent anatomical landmarks in the structural image upon which their activation map is overlaid.⁵¹² This is the approach that de Lange, Roelofs, and Toni deployed in their studies of hysterical hand paralysis. Yet, manual attribution of anatomical locations has one caveat. Researchers must be skilled enough to visually recognise anatomical structures that are characterised by considerable inter-subject variability. To acquire the requisite visual literacy, researchers are advised to practise working with anatomical images and thus “build up a 3D internal mental model of

511 Poldrack, Mumford, and Nichols, *Handbook*, 179. See also *ibid.*, 176; and Devlin and Poldrack, “Tedious Anatomy,” 1037.

512 Devlin and Poldrack, “Tedious Anatomy,” 1037; and Poldrack, Mumford, and Nichols, *Handbook*, 176.

neuroanatomy.”⁵¹³ But regardless of whether researchers prefer to rely on automated tools or to perform the anatomical attribution manually, even choosing which of the available brain atlases to use as a reference is a decision with significant epistemic consequences.⁵¹⁴

On the whole, my analysis has shown that determining the anatomical locations of statistically significant activations in an fMRI map is by no means straightforward but entails instead a step-by-step ‘reading’ procedure. While performing this procedure, researchers are required to continually make visual judgments about the functional maps by bringing them in relation to different types of images that visualise brain anatomy. Thus, the anatomical legibility of statistical brain maps depends on the researchers’ ability to embed these maps into a framework of intramedial references.⁵¹⁵ As we have seen, this framework consists of both the experimental subjects’ own structural imaging data and standardised images stemming from anatomical atlases.

Based on the analysis above, I argue that although a functional map is constructed as legible through statistical analysis that isolates the information of interest and makes it potentially accessible, the act of reading the map is entirely predicated on the combined use of multiple visualisations. In other words, researchers must actively engage with different, mutually complementary visualisations to visually articulate and thus finally gain access to the informational content of the statistical brain map. My analysis has foregrounded that researchers do not use visualisations as passive illustrations of an fMRI brain map. Instead, researchers deploy visualisations as flexible tools with which they perform a wide variety of operations. These operations include obtaining an efficient visual summary of the results, examining the shape and the spatial distributions of active clusters, as well as navigating around the visualised brain to inspect it across different anatomical locations and from multiple perspectives. Finally, I have shown that different types of visualisations facilitate the construction of the anatomical legibility of functional brain maps by bringing them into visual relations to other images. Therefore, the limits to what can be made visually distinguishable in a functional map during the process of result assessment fundamentally determine which aspects of the map can be made legible and thus comprehensible. In my opinion, it is this act of visual interpretation that, in the final instance, constitutes an fMRI map as a full-fledged indexical sign of the experimentally isolated pattern of brain activity.

513 Devlin and Poldrack, “Tedious Anatomy,” 1037. See also Jenkinson and Chappell, *Neuroimaging Analysis*, 193.

514 For a succinct overview of currently available types of atlases, including the so-called probabilistic atlases that are based on statistically-weighted composites of many individual brains, see, e.g., van Essen, “Windows on the Brain.”

515 Jäger, “Transcriptivity Matters,” 53.

3.5.2 Creating Publication Figures as Communication Tools

The previous section has highlighted how the flexibility with which functional maps can be visualised plays a crucial role during the working process, allowing researchers to actively explore their experimental results. The current section focuses on a distinctly different role accorded to visualisations that researchers specifically create for publishing their fMRI findings in scientific journals. In what follows, I will argue that publication images—which in research articles are designated as ‘figures’—are used as highly effective communication tools that visually convey and frame the experimental results in a particular way. Moreover, by returning to the de Lange, Roelofs, and Toni study, I will show that researchers must do two things to create figures that successfully perform this function. First, they have to construct their figures as multimodal visualisations that contain “words, numbers, and pictures.”⁵¹⁶ Second, they have to anchor the resulting figures into a specifically structured text of the research article.

As discussed earlier, in their initial fMRI study of conversion paralysis, de Lange, Roelofs, and Toni tested four different contrasts at the group level. Two of these group-level contrasts—the comparison between the affected and the unaffected, as well as between the left and the right hand—could be tested in two different directions. Hence, the researchers computed six statistical activation maps altogether.⁵¹⁷ In two of the maps thus obtained, no statistically significant activations were visible after thresholding.⁵¹⁸ Accordingly, the empty maps were not included in the publications. The published article, therefore, contains four visualised activation maps that the researchers organised into three separate figures.⁵¹⁹ In the subsequent section, we will discuss how the researchers interpreted these maps. But in what follows, we will first examine the structure of the figures with which de Lange, Roelofs, and Toni chose to communicate their empirical results to the scientific community.

What catches the eye even upon a cursory examination of the three figures is their distinctly composite character—multiple visualisations are unified under a joint caption. All three figures share the same visual organisation (fig. 3.14).⁵²⁰ On the left-

⁵¹⁶ Tufte, *Visual Display*, 10.

⁵¹⁷ The two contrasts were bidirectional. By subtracting the activations induced by the drawings showing the unaffected hand from the activations induced by the drawings of the affected hand, the researchers were able to identify the brain regions differentially activated by the affected hand. By reversing the direction of the subtraction, the researchers computed an additional map that isolated the differential activations specific to the unaffected hand. The same principle of directionality informed the comparison between the left and the right hand. The other two contrasts—the parametrised rotation-related increase in the activity versus baseline, and the interaction between the rotation-related differences and the hand affectedness—were not directional. See de Lange, Roelofs, and Toni, “Self-Monitoring,” 2054–55.

⁵¹⁸ The researchers found no statistically significant activations for the interaction between the rotation-related differences and hand affectedness. They also found no activation for the healthy relative to the paralysed hand. De Lange, Roelofs, and Toni, 2054.

⁵¹⁹ The two maps created by analysing two different directions of the bidirectional contrast between the left and the right hand were joined into a single figure. See de Lange, Roelofs, and Toni, 2056.

⁵²⁰ For this reason, I chose to reproduce here only one of these three figures.

hand side of each figure are the anatomical visualisations of the statistical map for the given contrast. These visualisations display a grey-scale structural slice of a brain encased inside a skull. The structural slice is overlaid with red-to-yellow colour-coded clusters of voxels that have been declared active in relation to the contrast specified in the respective caption. The orientations of the slices vary across figures, showing the brain either in the transversal, coronal, or sagittal cross-sectional plane. In the upper left corner, each slice is labelled with a single coordinate that specifies the location of the image plane within the standard space.

To the right of each anatomical visualisation is a bar graph that displays the estimated effect sizes. Each bar in the respective graph denotes the averaged strength of the BOLD response induced by the experimental conditions entailed in the respective contrast (fig. 3.14, b). The captions clarify that each graph shows the estimated effect sizes for the activation cluster, whose anatomical location is highlighted in the neighbouring anatomical visualisation with a yellow dotted circle. The captions also state the level of significance at which the visualised maps were thresholded and designate the anatomical regions in which the visualised activations are located. Furthermore, the captions refer the reader to two separate tables that entail the standard space coordinates of the peak activations. The stand-alone tables contain additional quantitative details, such as the standard space coordinates of all activated clusters, the sizes, and statistical values of the clusters, as well as their corresponding anatomical labels (fig. 3.15).

Several aspects are worth noting concerning the above examples because they are representative of how fMRI studies use publication figures to convey their results.⁵²¹ Most significantly, although the types and number of visualisations may vary across studies, the figures commonly comprise diverse visual components united under a joint caption. These components mutually complement one another, while each fulfils a specific function. For example, the purpose of the anatomical visualisation—which in all fMRI publications is the central and indispensable component of the figure—is to allow a clear localisation of the activations. Hence, researchers are guided primarily by pragmatic goals when deciding whether to overlay their statistical map onto a single or onto multiple structural 2D slices or, alternatively, to use a 3D brain rendering instead. Their professed concern is to impart maximum visibility to the locations of the activated clusters by displaying them on structural images in which the salient anatomical landmarks are easily identifiable.⁵²²

Notwithstanding the crucial role of anatomical visualisations in effectively transmitting the informational content of fMRI maps, I nevertheless want to suggest that this particular visual format has an additional rhetorical function. What I mean is that the images' "homogenous graphical language" facilitates the framing of the study's empirical results as straightforward, clear-cut, and unambiguous.⁵²³ The anatomical visualisations materialise only the polished outcome of a long and convoluted chain of

521 For examples of similar visualisation strategies, see, e.g., Aybek et al., "Life Events"; Morris et al., "Avoidance"; and Voon et al., "Involuntary Nature."

522 Huettel, Song, and McCarthy, *Imaging*, 370–71.

523 Latour, *Pandora's Hope*, 66.

transformations. At the same time, they hide the multitude of interpretational choices that went into producing the resulting statistical map. In doing so, the seemingly clear-cut anatomical visualisations suppress alternative interpretations that could have been produced from the same fMRI dataset had the researchers made different choices.

Figure 3.14. Visualisations of the statistical activation map for the contrast between the affected and the unaffected hand. (a), (c) and (e): anatomical visualisations; (b), (d), and (f): estimated effect sizes for select activation clusters. From: de Lange, Roelofs, and Toni, "Self-Monitoring," 2056, fig. 3. ©Elsevier.




Fig. 3. Regions showing higher activity for the affected than the unaffected hand. Anatomical localization and effect sizes ($\pm$ S.E.M.) of clusters showing overall (i.e., not rotation-related) higher activity for the affected hand than for the unaffected hand. There was higher activity for the affected limb in the left superior temporal cortex (a and b), medial prefrontal cortex (c and d), and the right superior temporal cortex (e and f). Exact stereotactic coordinates are given in Table 3. Other conventions as in Fig. 2.

A pertinently created anatomical visualisation provides the expert reader with an easily readable overview of the spatial distributions of the statistically significant activations computed for a particular contrast between experimental conditions. The caveat, however, is that this type of visualisation cannot communicate the precise quantitative information about the statistical significance of the anatomically displayed results. Since detailed statistical information is crucial for the informed reader in the scientific context, the publication figures entail additional visual elements. Such additional visual elements are typically derived from the visualisations researchers used

during the process of evaluating their statistical maps. Hence, as in the example above, the publication figures often combine anatomical visualisations with bar graphs (see fig. 3.14). Such graphs display in a visually straightforward manner the quantitative information about the estimated relative strengths of the neural responses induced by the mutually contrasted experimental conditions in an activation cluster of interest. The graphs thus allow the expert reader to assess the quality of the underlying data that went into producing the resulting statistical maps.

Figure 3.15. Table listing the statistical values computed for the contrast between the affected and the unaffected hand and the bidirectional contrast between the left and the right hand. From: de Lange, Roelofs, and Toni, “Self-Monitoring,” 2055, table 3. ©Elsevier.

Table 3 Cerebral data—activation differences							
Contrast	Region	Pseudo- <i>T</i> value	Cluster size	Corrected <i>p</i> -value	Stereotactic coordinates		
					<i>x</i>	<i>y</i>	<i>z</i>
Affected > unaffected	Medial frontal cortex	5.5	1303	0.035	8	44	−24
		5.2			−12	62	32
		6.2			−36	48	34
	Parietal operculum (PO4)	5.8	1065	0.039	−58	−6	10
	Superior temporal sulcus	5.1			−52	−36	−4
	Superior temporal gyrus	5.9			68	−28	10
Left hand > right hand	Primary motor cortex	5.4	4673	0.0039	16	−40	70
	Precentral gyrus	7.0			32	−10	68
Right hand > left hand	Primary motor cortex	7.1	1525	0.0098	−6	−36	64
All reported coordinates are in MNI (Montreal Neurological Institute) space. Stereotactic coordinates denote the peak of the clusters surviving correction for multiple comparisons.							

Yet, even such multimodal figures do not suffice to transport the requisite information with adequate precision. Therefore, in addition to the figures, almost all published studies provide stand-alone tables that are visualised separately under their own captions (see fig. 3.15).⁵²⁴ As in the example above, such tables summarise multiple quantitative aspects of the statistical map, which the anatomical visualisations are unable to convey. Moreover, it appears to me that the tables also fulfil a rhetorical function within the published article. Specifically, I suggest that the statistical tables add to the persuasiveness of the results presented in the anatomical visualisations of the maps. They do so by providing a strictly quantitative perspective, thus linking the results to the initial measurement.

In effect, the presentation of the empirical results in the research article is partitioned into two panels with different characteristics. One panel has a predominantly visual character. This panel entails a composite figure that displays anatomical visualisations of the statistical brain map and, in many cases, an accompanying bar chart of the estimated effects sizes. The other panel comprises exclusively numerical elements, as it contains the table with various statistical values and a list of the coordinates. On their own, both the visual and the numerical modalities provide only a partial insight into the results. It is only by bringing these mutually complementary elements in relation to one another that the expert reader can

524 Poldrack, Mumford, and Nichols, *Handbook*, 179.

reconstruct the ‘full picture.’ Therefore, I argue that the visual and numerical aspects of fMRI maps are semantically enmeshed not only during the process of working with the data but also when viewing the results within the published research article.

Interestingly, in many articles—our case study included—the figures and tables that refer to the same statistical maps are often presented on different pages. Hence, the reader has to switch back and forth between the figures and maps to successfully decipher the visual information that is spread across the article. Although I am speculating here, it is almost as if the implicit aim of such a layout is to emphasise the inability of any single modality to convey the full complexity of the neuroimaging findings. But, regardless of a particular layout, the fact remains that to understand the findings of a neuroimaging study in their full complexity, the skilled reader of a published article has to do much more than simply glance at a single anatomical visualisation of a statistical brain map. To grasp the results, the reader is forced to emulate the researchers’ working process, during which they continually juggle not only the numerical and visual aspects of the data but also rely on multiple modes of visualisation.

However, there is one crucial difference between, on the one hand, the process of viewing the fMRI figures published in a research article and, on the other hand, working with visualisation during the assessment of fMRI results. As analysed in the previous section, throughout the working process, statistical maps remain firmly embedded into the chain of inscriptions that produced them. We have seen that digital interfaces and software packages allow a fluid movement along this chain by making all the previous inscriptions accessible at a click of a button. Just as importantly, the digital interfaces also enable researchers to use the visualisations in highly dynamic and interactive ways as tools for actively exploring the data. By contrast, in the published form, this circulating flow of the mutually interconnected inscriptions is arrested and displaced by a limited set of fixed figures and tables that clearly and persuasively display the results in their distilled form. Unlike the malleable visualisations used during the working process, the images in publications are no longer interactive or ‘surfable.’ Instead, to use the term introduced by Latour, the published figures are specifically designed to function as immutable mobiles.⁵²⁵ In short, these figures are easily reproducible inscriptions that enable the displacement of information without any further transformation.

The unavoidable downside of the newly won immutability of these images is that the process of their visual fixation effectively cuts them off from the chain of references that produced them. What remains invisible and illegible in the published images—even for an expert—are the exact details of the interpretational choices, theoretical and practical assumptions, modelling approximations, filtering and standardisation that researchers have undertaken to arrive at the result visualised in the published figures and tables. But, if fully isolated from the chain of transformations through which it was produced as an indexical sign, an fMRI map would become meaningless in the scientific context. In such a case, researchers could no longer use an fMRI map to make

525 Latour, “Visualization and Cognition,” 19–22.

judgments about the putative cerebral dysfunctions that give rise to hysterical paralysis or, more generally, any other phenomenon under study.

In fact, to enable their published fMRI figures to retain their referential relation to hysteria patients' active brains, researchers have to anchor these figures into a specifically structured text of the research article. For this reason, the major portion of an fMRI-based research article entails a detailed narrative reconstruction of the original chain of references through which a particular statistical brain map was created. In other words, the lack of the physical—i.e., digital—access to the cascade of the previous inscriptions is thus substituted by a step-by-step description of the operations that went into producing the visualised statistical maps. Such narrative descriptions include the criteria of the participant selection, details of the task design and its concrete implementation, the parameters of the data acquisition and preprocessing, as well as the researchers' decisions and mathematical operations that shaped various stages of the statistical analysis. Without such a sufficiently precise narrative reconstruction of their underlying chain of references in the main text of the research paper, the fMRI maps are unable to refer indexically to the measured brains and, as a result, lack the epistemic efficacy.⁵²⁶ With the aim of preventing such situations, there have been repeated calls in the neuroscientific community to establish standardised guidelines for reporting the results of fMRI-based research.⁵²⁷

In sum, despite its apparent ability to summarise complex results and endow them with visibility, legibility, and materiality, a composite visualisation of a statistical map within a published article is “a strange transversal object, an alignment operator.”⁵²⁸ The evidentiary status of such a figure arises from a complex interplay of its heterogeneous visual and numerical components, the accompanying caption and the main text of the research article. Therefore, I argue that it is not the function of the fMRI maps to illustrate the text of the published research article. Instead, as foregrounded by my analysis, the major portion of the research article has an auxiliary, descriptive character that serves to validate the fMRI maps by reconstructing the referential chains that underpinned the maps' production. In short, not the text but the images are of central importance in an fMRI paper since they present the paper's empirical findings. However, as we are about to see in what follows, the role of the text shifts considerably in the final sections of the research article. That is, in the article's section referred to as the ‘discussion,’ the text acquires a more active role in constructing the meaning of the statistical map. Let us now take a close look at how and why such a shift occurs.

3.5.3 Staging the Meaning of Functional Brain Maps

Until now, I have delineated the cascade of operations through which fMRI maps are produced and how the thus isolated patterns of task-induced brain activities are

526 See Poldrack et al., “Guidelines for Reporting,” 409–14.

527 Poldrack et al., 409–14.

528 Latour, *Pandora's Hope*, 67.

visualised in the form of fixed composite figures. Yet, even if the figures are adequately designed to convey the empirical results clearly, the meaning of the anatomically circumscribed activations they display is not self-evident. Hence, in the final stage of an fMRI study of a hysterical symptom, researchers have to posit an interpretation of the experimentally isolated brain activities in terms of the symptom's potential neural mechanism. To do this, researchers have to address the following questions: Which aspects of the activation patterns visualised in the patients' functional maps are aberrant? Which distinct cognitive functions are associated with these aberrant patterns of brain activity, and how do they give rise to the hysterical symptom under investigation? Deploying Ludwig Jäger's concept of transcriptivity,⁵²⁹ I will argue that, by answering these questions, researchers stage the symbolic meaning of their fMRI maps. They do so by inscribing each map into a specifically constructed frame of intramedial and intermedial references. In this section, I will examine this process by drawing on the example of the two mutually related de Lange, Roelofs, and Toni studies of hysterical hand paralysis.

In the main text of their initial study, de Lange, Roelofs, and Toni listed somewhat cursorily all the steps that went into producing their fMRI maps before moving to the description of their empirical results.⁵³⁰ First, they summarised the patients' behavioural measurements. These included the subjects' reaction times and their respective task performance error rates. Next, the researchers delineated their fMRI results by stating the anatomical locations of the task-induced neural activations obtained for different contrasts of the experimental conditions. In the final section of the article, de Lange, Roelofs, and Toni finally turn to developing an overarching narrative interpretation of their results in a step-by-step procedure. Yet crucially, it was based on this narrative interpretation of the fMRI maps that de Lange, Roelofs, and Toni were able to suggest a possible functional mechanism underlying the loss of volitional movement in conversion paralysis. Thus, in what follows, we need to analyse how the researchers constructed their interpretation of the maps.

De Lange, Roelofs, and Toni began their interpretation by focusing on the fMRI map that disclosed in which brain regions the neural activity intensified in response to the increasing biomechanical complexity of the task for both the affected and the unaffected hand. The map showed that the increasing rotation level of the implicitly imagined movements induced an equivalent pattern of neural activity for both hands.⁵³¹ As the researchers explained, the resulting pattern of activations was located in the dorsal parietal and premotor cortex of the patients' brains—the areas known to play crucial roles in planning voluntary movements.⁵³² De Lange, Roelofs, and Toni then compared this map with the activation maps generated by previous fMRI studies that had been conducted on healthy individuals. Some of the previous studies used a similar implicit

529 Jäger, "Transcriptivity Matters," 49–50.

530 The authors described the participant selection criteria, the task, the parameters of the acquisition, preprocessing steps, and the basic aspects of their statistical analysis. See de Lange, Roelofs, and Toni, "Self-Monitoring," 2052–53.

531 See de Lange, Roelofs, and Toni, 2054, fig. 2.

532 De Lange, Roelofs, and Toni, 2055–56.

motor imagery task to induce imagined movement in experimental subjects, whereas others investigated the initial phase of an actually performed hand motion.⁵³³ Based on the similarity between the anatomical locations of the activated clusters across the respective maps, de Lange, Roelofs, and Toni concluded that their patients exhibited a normal activation pattern in the motor cortex. This example clearly demonstrates that a decision on whether an experimentally isolated pattern of brain activity can be categorised as ‘normal’ rests on a comparison with an already established pattern of ‘normal activity.’

In the next step, de Lange, Roelofs, and Toni mobilised two additional empirical findings to reinforce the claim that their patients showed normal motor cortex activation when implicitly imagining movement. First, the researchers emphasised that the behavioural data showed no statistically significant differences in the patients’ task performances between the paralysed and the healthy hand, either regarding reaction times or error rates.⁵³⁴ Additionally, de Lange, Roelofs, and Toni pointed out that the hypothesis testing of the contrast that compared rotation-related differences between the affected and the unaffected hand resulted in an empty map. They interpreted this lack of activation as further evidence supporting the claim that the task’s increasing biomechanical complexity induced comparable cerebral responses for both hands.⁵³⁵ Based on these converging results, de Lange, Roelofs, and Toni conjectured that individuals with conversion paralysis “can readily imagine actions of both their unaffected and affected hand, using the same cerebral resources as healthy participants.”⁵³⁶

Two aspects of this conjecture are significant. At this point, the researchers already started generalising their findings beyond their sample of participants to conversion disorder (i.e., hysteria) patients in general. Moreover, their assertion also provided an a posteriori validation of the adequacy of their experimental task for isolating the neural activity specific to the loss of volitional movement in conversion paralysis. Specifically, by combining fMRI and behavioural data, de Lange, Roelofs, and Toni have proven that their patients were able to carry out the experimental task equally well with both hands. Drawing on this proof, the researchers could, in turn, claim that another fMRI map, which was calculated from the same dataset and displayed the differential neural activity between the affected and the unaffected hand, was not confounded by a potential difference in the task performance.⁵³⁷ Hence, de Lange, Roelofs, and Toni constructed the meaning of the individual maps not only in relation to the patterns of ‘normal’ brain activity provided by other studies but also by cross-referencing different empirical findings within their own study. In each case, the ascription of meaning was distinctly relational as it entailed a comparison of the map in question either to other fMRI maps or to another type of data—i.e., error rates and reaction times.

533 For the list of these studies, see de Lange, Roelofs, and Toni, 2056.

534 De Lange, Roelofs, and Toni, 2054.

535 De Lange, Roelofs, and Toni, 2056.

536 De Lange, Roelofs, and Toni, 2056.

537 De Lange, Roelofs, and Toni, 2055.

The same strategy informed the researchers' interpretation of the two maps generated by the bidirectional contrast between the imagery of the left and the right hand. First, by referencing the findings of previous fMRI studies, de Lange, Roelofs, and Toni established that the differential activations for both directions (i.e. left to right and right to left) exhibited normal patterns.⁵³⁸ The researchers then turned to cross-referencing the maps within their study by comparing the two maps that contrasted the left and the right hand with the map that displayed the differential activation between the imagery of the affected and the unaffected hand. This comparison revealed that there were no overlapping activation patterns across the maps. De Lange, Roelofs, and Toni thus concluded that the same neural processing underpinned both the left- and right-hand conversion paralyses.⁵³⁹ The implication entailed in this conclusion was that the differences in the laterality of paralysis across patients did not confound the map computed by contrasting the imagery of the affected and the unaffected hand. It is worth noting that in this particular case, not the similarities but the differences across the maps proved to be of semantic relevance to the interpretation.

So far, we have seen how de Lange, Roelofs, and Toni mobilised the behavioural data and multiple fMRI maps to gradually develop their claim that the map displaying the differential brain activity between the motor imagery of the affected and the unaffected hand can indeed provide insights into the putative neural mechanism underpinning conversion paralysis. It is only at the end of this process of semantic contextualisation that the researchers finally turned to revealing this mechanism. But to do this, they first had to perform an additional semantic operation.

Known in the neuroscientific context as "reverse inference," this semantic operation entails reasoning backwards from the activity of a particular brain region to a specific cognitive process.⁵⁴⁰ As I am about to show, this kind of non-statistical inference involves the ascription of meaning that is extraneous both to the visual content of the fMRI maps and the experimental setup that generated them. Instead, this kind of non-statistical inference relies exclusively on textual—i.e., intramedial—references to other fMRI studies that have postulated a putative link between a brain region of interest and a cognitive function. Importantly, the major caveat of this approach is that neuroscientific research, on the whole, has not yet provided evidence of any one-to-one mapping between brain anatomy and function.⁵⁴¹ Consequently, the activation of any single region can be attributed to different cognitive processes. By analysing how de Lange, Roelofs, and Toni approached this problem, I argue that the critical step in instituting a statistical activation map as a symbolic sign of a particular cognitive process consists in the semantic staging of selective references to other fMRI studies.

First, de Lange, Roelofs, and Toni listed multiple brain areas that showed greater activation during the implicit imagery of the affected as opposed to the unaffected hand (see fig. 3.14). These included the left and right superior temporal cortex and

538 De Lange, Roelofs, and Toni, 2055.

539 De Lange, Roelofs, and Toni, 2055.

540 A more common type of non-statistical inference in neuroimaging runs in the opposite direction: "if cognitive process X is engaged, then brain area Z is active." Poldrack, "Cognitive Processes," 59.

541 Poldrack, 60–61.

three different regions within the prefrontal cortex. For reasons they did not disclose, de Lange, Roelofs, and Toni chose to focus primarily on one of these regions—the ventromedial prefrontal cortex (vmPFC) (fig. 3.14, c). Presumably, their choice was motivated by the fact that the vmPFC had already been implicated in two previous neuroimaging studies of hysterical paralysis published in 1997 and 2000.⁵⁴² The authors of the previous studies of hysterical paralysis had postulated that the pathologically increased activity in this particular area of the prefrontal cortex was associated with the functional disturbance of the cognitive process called inhibitory control. Under normal conditions, the purpose of inhibitory control is to suppress the execution “of inappropriate motor responses.”⁵⁴³ But Marshall et al. argued that in hysteria patients, the pathological activation of the prefrontal cortex led to “unconscious inhibition” of the normal activity in the motor cortex, thus resulting in the abolishment of volitional movement in the patients’ affected limbs.⁵⁴⁴ Interestingly, as explicitly stated by Marshall et al., their interpretation partly overlapped with the neurophysiological mechanism Charcot had posited more than a century earlier as the potential basis of hysterical paralysis.⁵⁴⁵

However, the interpretation that posited increased inhibitory control of the motor system ran contrary to one of the fMRI findings that de Lange, Roelofs, and Toni made in their study. As discussed above, one of their fMRI maps showed that conversion/hysteria patients activated the same motor-related brain structures as healthy subjects during imagined movements of the paralysed hand. Drawing on this map, de Lange, Roelofs, and Toni contradicted the reverse inference suggested by Marshall et al. Instead, they posited an alternative interpretation by contextualising

542 De Lange, Roelofs, and Toni quoted two studies: Marshall et al., “Hysterical Paralysis”; and Halligan et al., “Hypnotic Paralysis.” However, there were inconsistencies across the three studies concerning the standard space coordinates of the peak activations they identified. Moreover, the studies also used different anatomical labels to designate the activated areas in the prefrontal cortex. Nevertheless, both Halligan et al. and de Lange, Roelofs, and Toni explicitly claimed that their results mutually overlapped in terms of spatial distribution. Compare Marshall et al., “Hysterical Paralysis,” B3–6; Halligan et al., “Hypnotic Paralysis,” 987; and de Lange, Roelofs, and Toni, “Self-Monitoring,” 2056. It should also be noted that, unlike de Lange, Roelofs, and Toni, the two earlier studies used PET and not fMRI. As discussed in chapter 2, PET has considerably lower spatial resolution than fMRI and, therefore, results in a less precise localisation of neural activity.

543 Marshall et al., “Hysterical Paralysis,” B2.

544 Marshall et al., B6.

545 Marshall et al., B5–6. Charcot and Halligan et al. had in common the conjecture that the unconscious inhibition led to a suppression of the activity in the motor cortex, which, in turn, caused hysterical paralysis. Yet, Charcot localised the inhibition in the cerebral motor centres. Conversely, Halligan et al. implicated additional cortical areas such as the vmPFC and the anterior cingulate cortex, thus suggesting a considerably more complex mechanism involving interactions across multiple brain regions. Furthermore, Halligan et al. drew their conjecture based on the functional brain map that displayed their single patient’s brain activations. Charcot, instead, relied on far more indirect images that visualised the spatial patterns of the paralysed patients’ accompanying anaesthesia. For a discussion of Charcot’s views on the neural basis of hysterical paralysis and how he developed them using hypnotic experiments and visualisation techniques, see section 1.3.2.

their activation map within the cognitive framework drawn from the paradigm of resting-state fMRI research.

According to multiple resting-state studies, the vmPFC is part of the so-called default-mode network. This network entails multiple, mutually interacting brain areas, whose activity is high during periods of wakeful rest, when a subject is engaged in self-referential cognitive activities, such as monitoring one's own mental states.⁵⁴⁶ Conversely, the activity of the default-mode network decreases as soon as the subject is engaged in the execution of external, goal-oriented tasks that require "an attenuation of self-focused attention."⁵⁴⁷ Based on these findings, de Lange, Roelofs, and Toni suggested that the increased activity of the vmPCF displayed by their fMRI map arose from hysteria patients' inability to deactivate this region while imagining movements of the affected hand. Specifically, in their patients, the activity of the vmPCF remained "at resting-state levels" even during the task execution.⁵⁴⁸ De Lange, Roelofs, and Toni attributed this aberrant activity of the vmPCF to hysteria patients' abnormally increased self-monitoring processes. In short, the researchers concluded that, unlike healthy subjects, hysteria patients could not attenuate their self-referential mental activity when they were engaged in goal-directed tasks.

To further substantiate their reverse inference, de Lange, Roelofs, and Toni drew attention to the aberrant activity of two other clusters in their fMRI map—the left and right superior temporal cortex (see fig. 3.14, a and e). The researchers suggested that the abnormally increased activity in these two regions during imagined movements of the paralysed hand potentially reflected "heightened monitoring of actions with the affected limb, but in the visual domain."⁵⁴⁹ De Lange, Roelofs, and Toni justified this claim by referencing findings from neurocognitive research into the functions of these two regions.

It is important to emphasise that de Lange, Roelofs, and Toni did not empirically refute the interpretations of the previous neuroimaging studies, which had attributed the formation of hysterical paralysis to higher-order inhibitory processes. Rather, they diverged from the authors of the previous studies by choosing a different set of references on which they based their reverse inference. Although their activation map partly replicated the findings of Marshall et al., through this act of referential re-staging, de Lange, Roelofs, and Toni shifted the interpretation of the neural activity detected into a new semantic context. By attributing the increased activity of the same prefrontal brain region to a different cognitive function, de Lange, Roelofs, and Toni effectively silenced the competing accounts.⁵⁵⁰ As a result, and at least for the time being, this particular neuroanatomical region ceased to function as a symbolic sign of heightened inhibition of motor processes in hysterical paralysis. Instead, de Lange,

546 Gusnard and Raichle, "Baseline," 691–92.

547 Gusnard and Raichle, 692.

548 De Lange, Roelofs, and Toni, "Self-Monitoring," 2056.

549 See de Lange, Roelofs, and Toni, 2057.

550 I am using the term 'silencing' here in Jäger's sense. Jäger has argued that each semantic transcription suppresses and thus silences alternative meanings that had been generated by "different transcriptive situations." Jäger, "Transcriptivity Matters," 62.

Roelofs, and Toni instituted the increased activity of the vmPFC as a sign of hysteria patients' pathologically altered self-focused monitoring. But, there were two caveats.

First, the newly assigned meaning remained somewhat vague. Multiple functional neuroimaging studies had suggested that the vmPFC might contribute to the integration "of continuous cognitive and emotional processes" through "online monitoring of associations between sensory information, responses and outcomes under changing circumstances."⁵⁵¹ However, de Lange, Roelofs, and Toni had to concede that, based on their study design, it was impossible to determine if the hysteria patients' increased self-monitoring focused on the potential sensorimotor or emotional consequences that arose from imagined movements of the paralysed hand.⁵⁵² Hence, in their study, the exact nature of the hysteria patients' purported self-monitoring processes remained undetermined.

Second, as pointed out by Jäger, any discursive ascription of meaning is inherently unstable. Its semantic legitimacy is grounded in the symbolic act that postulates its own interpretation not only as preferred but possibly also as the only correct interpretation.⁵⁵³ At the same time, each ascription of meaning also necessarily "opens the realm for competing transcriptions," thus setting in motion "the iterative-endless game of semantic re-staging."⁵⁵⁴ The fragility of discursively instituted meanings is particularly pronounced in highly dynamic research contexts—such as cognitive neuroscience, in general, and fMRI-based investigation of hysteria, in particular. In such contexts, each new study tends to recast the conclusions derived from previous experimental findings. As we will see in the following chapter, subsequent fMRI studies of hysterical paralysis continued to readdress the potential role of both inhibition and increased self-monitoring as possible mechanisms underpinning this symptom. Yet such semantic re-staging is not only limited to interactions between different, mutually competing researchers teams. Instead, the semantic transcription can also be undertaken by the very same researcher team.

A pertinent example of the latter is provided by the subsequent fMRI study of hysterical paralysis by de Lange, Roelofs, and Toni, in which these researchers used the same dataset to compute statistical connectivity maps.⁵⁵⁵ Three years after their initial study, de Lange, Roelofs, and Toni decided to address the question that had, in the meantime, arisen in fMRI-based hysteria research and to which their suggested neural mechanisms of increased self-monitoring could not provide an adequate answer. The question was: How is the increased prefrontal activity related to the consistently reduced responses in motor and sensorimotor brain areas?⁵⁵⁶

To answer this question, de Lange, Roelofs, and Toni used the PPI analysis.⁵⁵⁷ They aimed to determine how the functional connectivity between the three prefrontal

551 Gusnard and Raichle, "Baseline," 692.

552 De Lange, Roelofs, and Toni, "Self-Monitoring," 2057.

553 Jäger, "Transcriptivity Matters," 64.

554 Jäger, 64.

555 See de Lange, Toni, and Roelofs, "Altered Connectivity."

556 De Lange, Toni, and Roelofs, 1783.

557 For a detailed discussion of the analysis pipeline through which de Lange, Toni, and Roelofs computed their connectivity maps, see section 3.4.4.

areas—which they had isolated in their previous study—and the rest of the brain changed depending on whether hysteria patients imagined movements of their affected or unaffected hand. The computed connectivity maps disclosed that the vmPFC did not exhibit statistically significant coupling with any parts of the brain's sensorimotor network. This was inconvenient because the vmPFC was the region to which de Lange, Roelofs, and Toni previously accorded the central role in their proposed neural mechanism of conversion paralysis. Due to this new finding, the researchers were forced to concede that the “vmPFC does not directly impinge on the sensorimotor system.”⁵⁵⁸

Yet, fortuitously, another frontal area called the dorsolateral prefrontal cortex (dlPFC)—whose coordinates were listed in their previous study among the activation peaks—showed strong connectivity with several sensorimotor regions. Specifically, when viewing the motor imagery of the affected versus the unaffected hand, patients showed aberrantly increased positive coupling between the dlPFC and the premotor cortex.⁵⁵⁹ The same contrast of the experimental conditions also induced a negative coupling between the dlPFC and the primary sensorimotor cortex.⁵⁶⁰ Referencing multiple neuroimaging studies, de Lange, Roelofs, and Toni conjectured that these abnormal connectivity patterns reflected hysteria patients' functional disturbance in the “formation of action plans of the affected arm.”⁵⁶¹ They further concluded that this particular disturbance could be implicated in the loss of volitional movement in hysterical paralysis.

In effect, by computing a different type of map from the same fMRI dataset, de Lange, Roelofs, and Toni obtained additional empirical findings that did not seamlessly fit into their previously proposed neural mechanism. The researchers thus had to narratively re-stage their initial interpretation while, at the same time, trying to preserve its legitimacy. In the end, they were unable to reconcile the old and new findings into a single, internally consistent narrative. They settled instead on a slightly less elegant solution. According to their updated interpretation, two disparate neural mechanisms—heightened self-monitoring and a disturbance in action selection—played roles in conversion/hysterical paralysis.⁵⁶²

As we have seen, de Lange, Roelofs, and Toni attributed each of these two disparate mechanisms to different anatomical parts of the prefrontal cortex. Problematically, this updated interpretation could not explain how the two mutually disparate presumed mechanisms interacted with each other to give rise to the loss of volitional movement in hysterical paralysis. It can thus be said that de Lange, Roelofs, and Toni failed to unambiguously identify the “objective neural correlates of functional mechanism” underpinning hysterical paralysis.⁵⁶³ Nevertheless, it should be emphasised that they made a significant contribution to the fMRI investigation of hysteria by opening up

558 De Lange, Toni, and Roelofs, “Altered Connectivity,” 1786.

559 De Lange, Toni, and Roelofs, 1785.

560 De Lange, Toni, and Roelofs, 1785–86.

561 De Lange, Toni, and Roelofs, 1785–86.

562 De Lange, Toni, and Roelofs, 1786.

563 De Lange, Toni, and Roelofs, 1782.

new lines of research. Far from merely replicating Charcot's initial conjecture that hysterical paralysis was due to the functional inhibition of the motor cortex, de Lange, Roelofs, and Toni innovatively proposed two entirely different and more complex neural mechanisms. The following chapter will show that their findings and hypotheses were taken up and further developed by subsequent fMRI studies.

In summary, my analysis has demonstrated that to examine the epistemic potential of fMRI in current hysteria research, we should neither approach it as a transparent window into the putative neural mechanisms of this disorder nor as a source of pretty but baseless pictures. Instead, in this chapter, I have argued that fMRI is better understood as an investigation tool whose deployment in contemporary hysteria research has opened up radically new possibilities for generating novel insights into this mysterious disorder. Yet, as highlighted by my analysis, the application of fMRI is also coupled with considerable methodological challenges.

In order to use fMRI in epistemically productive ways, researchers must properly align numerous material and discursive operations along a consistent chain of transformations, whose individual stages I have delineated in this chapter. The product of such an alignment—which entails automated algorithmic processes and active human judgments—is a set of functional brain maps. These maps are curious, hybrid objects that arise from a synthesis of measurement and modelling. Although, in essence, fMRI maps are mathematical entities, their informational content is accessible to human judgment through various forms of visualisation. It is through the combined use of multiple visualisations that researchers make sense of the fMRI maps and the underlying data, thus using them to produce new medical knowledge of hysteria. Therefore, I have argued that the limits of what can be visualised in fMRI maps determine the limits of their epistemic efficacy.

Moreover, we have seen that with each functional brain map, researchers actively create new phenomena that do not exist independently of the complex chain of medium-specific operations through which the maps were produced. These operations start with the initial inscription of physical traces into the fMRI data. They are then followed by a long series of transformations whose aim is to articulate the initial traces with sufficient clarity and precision. Provided that these operations were performed according to the current standards of the scientific community and aligned into an unbroken chain of references, the resulting maps are constructed as highly mediated indexical signs of the otherwise inaccessible neural activity of interest. Hence, to be epistemically relevant, the identified pattern of neural activity has to be grounded in the physical measurement of active brains. But just as importantly, this pattern of neural activity also has to be artificially created through operations of statistical averaging and standardising. As I have shown, these operations are necessary to isolate the activity of interest from incidental cerebral processes and to purify it from individual subjects' idiosyncrasies.

Yet, even after researchers have successfully performed all these time-consuming operations, they still face a crucial challenge. In the final step, researchers have to

provide a meaningful interpretation of their empirical findings in terms of related cognitive processes. Although fraught with difficulties, this step potentially carries the most significant epistemic impact because it allows researchers to use functional brain maps to produce shifts in how hysteria is conceptualised in the medical context.

For this reason, the following chapter will examine how fMRI-based hysteria research, on the whole, has begun to shift the medical understanding of the present-day manifestations of hysteria by producing new—although still tentative—empirical insights into the neural basis of this disorder. We will see that some of these new insights partly overlap with Charcot's long-challenged views on hysteria, whereas others open up entirely new epistemic perspectives on this still vaguely understood disorder.

