

What's so Good About Functional Images?

Jim Bogen

Pitzer College, Pittsburgh Center for
Philosophy of Science

i. We know experimental results have epistemic virtues. Otherwise, they'd be no good for developing and evaluating theories. We don't know enough about what these virtues are or where they come from—especially for cases like functional imaging of the brain (FI). To produce functional images, things are arranged so that radio or radiation signals are emitted from blood circulating through the brains of subjects at rest or engaged in mental tasks. Functional images are computed from, but are by no means intended to represent or report these signals. They are of interest only to the extent that they can be used to calculate blood volume or oxygenated hemoglobin levels indicative of cognitively significant levels of neuronal activity (CSNA). I'll call these biological indicators of CSNA 'BI'.

Neuroscientists do not develop and evaluate their theories by examining and reasoning from signal detections any more than physicians diagnose patients by reasoning from measurements of sound waves emitted by patients reporting their symptoms. The evidence neuroscientists use consists of images computed from signal detection records. The connections between radiation and radio signals and CSNA are indirect and imperfectly understood. Functional images are generated by computers running programs which embody statistical, biological, physical, and other assumptions, together with conventions, decisions, and stipulations developed to deal with technical problems as they come up. These are only partially constrained by what is known about the relevant physiological mechanisms and the production and detection of the signals.

This evidence is a far cry from observation reports of observers' sensory experiences, from perceptions of publicly accessible items, and from the observation reports 20th century empiricists believed were the epistemically significant fruits of observation and experimentation..(*cite Schlick, Russell, Neurath, Carnap). They are not records of instances or counter-instances of theoretical claims they are used to evaluate, or of predictions derived from those claims. Unlike typical measurement data, they are not required to be particularly accurate or precise with regard to the locations or levels of BI or CSNA of the brains they are derived from.¹ And their use to test and develop theories typically does not involve much in the way of corrections or quantitative estimates of their inaccuracies.

Deborah Mayo's excellent account of experimental reasoning suggests that the epistemic virtues of evidence often derives from statistical analyses of the noisy results of a real world

¹ *Note that this isn't peculiar to FI evidence.

experiment undertaken to figure out what data might have been produced from an ideal version of the experiment free from significant sources of error. Her account applies to cases where the experimenter can specify assumptions whose satisfaction would secure a satisfactorily error--free version of the experiment and can use error analysis to

... find a way (through corrections or whatever) of arriving at a desired data statement that effectively subtracts out (or renders harmless) any violation of the assumptions. (Mayo [1996], p.136).

Some of the statistical operations performed on signal detection records and estimates generated from them to compute a functional image (e.g., corrections for signal detection error) are meant to accomplish what Mayo has in mind. But it is implausible that all of them serve to bring imaged levels and locations of BI closer to what they would have been had the images been produced without significant error. One reason for this is that the final products typically exhibit the anatomy of imaginary and significantly atypical ideal brains. Furthermore, although computations are undertaken to avoid certain kinds errors in calculating BI levels, there is no quantitative assessment of the degree to which the results conform to real magnitudes. Finally, some of the statistical manipulations (including threshold setting and averaging discussed below) can be expected to decrease rather than increase accuracy. If the purpose of effectively subtracting out, or rendering harmless, violations of assumptions which define an ideal FI experiment was to bring the imaged BI or CSNA levels and locations closer to those of anatomically and functionally normal brains, there is no reason to think FI statistical analysis succeeds. If correcting for or mitigating error does not improve accuracy it's hard to see how it could account for the evidential value of the images.

Functional images differ from the data whose use to investigate items which are not themselves observable Jim Woodward and I have discussed. (*cite Bogen Woodward) Such data (e.g., thermometer or pressure gauge readings, photomicrographs, astrographs, bubble chamber photographs, drawings, seismic or EKG tracings) owe their epistemic value to having been produced in such a way that the item the investigators would like to detect or measure can exert a detectable causal influence on them. Such data owe their evidential value to features from which investigators can infer facts about the item of interest. To identify and draw conclusions from these telltale features, investigators reduce the data and employ extensive analysis to eliminate or mitigate obscurities due to the effects of extraneous causal influences on their production. Our account applies to cases in which what makes good evidence good is analogous to the clarity and volume of those parts of a radio broadcast which are relevant to an audience's interest. By that I mean the easier it is to discern the features of the data which are indicative of the item of interest, and to discriminate them from the effects of obscuring factors, the better the data. What sort of accuracy and precision contributes to the epistemic value of a parcel of data depends upon, and varies with, the questions the data is used to answer, and the

strategies employed to interpret them. That is why Galileo's moon drawings could be epistemically valuable despite gross inaccuracies in their depictions of lunar craters and mountains.

Unfortunately, reducing, analyzing, and interpreting the radio and radiation signals is so essential to the production of functional images that this story doesn't apply straightforwardly to them; they are not so much data as graphic representations of interpretations of data. But suppose we ignore this and apply what Woodward says about data in Woodward [2000] to a functional image **e**, used to argue for a claim, **h**, about the CSNA involved in the performance of a mental task. Then the epistemic value of **e** would depend upon the reliability of the process which produced **e**. Reliability would be a matter of whether the causal influence of CSNA of interest on the production of **e** was such that a relation of counterfactual dependency obtains between **h** and **e**. Ideally, it would be counterfactually the case that **e** has the features it has just in case **h** is true. Less ideally, **e** would be just as it is just in case **h** is highly probable, or more correct, or more probable, than the alternatives it is compared to. At the very least **e** shouldn't be as it is regardless of whether **h** is true, highly probable, etc. To evaluate the epistemic value of a functional image we would have to find out whether any such counterfactual dependencies obtain. This, according to Woodward, is to be ascertained from facts about the physical, and computational mechanisms involved in the production of **e**. The reliability of the image producing process is what we would need to determine in order . Although data are significantly different from functional images, and although Woodward's purposes did not require him to consider in detail how reliability is assessed, I think what I have to say about functional images agrees with Woodward [2000] in spirit, but differs with regard to how much an understanding of functional image production can contribute to its epistemic evaluation.

The reliability of an evidence producing process always depends upon what questions the evidence is applied to and how it is used. For example, ordinary light microscopes which can produce clear enough images to answer some questions about structures in very thin slices of prepared tissue, cannot provide well enough resolved images to answer similar questions about structures a few millimeters below the surface of a living specimen. Light scattered back from structures in front of, and behind such structures spoils the resolution by interfering with light scattered back from structures on the focal plane. In many cases, to image them with an ordinary light microscope, one must kill or alter the structures by cutting the specimen into slices so thin that they lie close to the surface. By contrast, optical coherence microscopy (OCM) non-invasively delivers well resolved images of structures 2 mm below the surface of living tissue. While ordinary light microscopic imaging can be highly reliable for such purposes as counting cells or organelles in a preparation, slicing and preparing the tissue prevents its repeated imaging of the same living specimen for such purposes as tracking changes in shapes, relative sizes, and

positions of cells in developing embryological tissue. Because it allows for this, OCM imaging is reliable for investigation to which ordinary light microscopy is unsuited. The limitations of OCM image precision and accuracy can be quantitatively determined by appeal to facts about the process by which they are produced without reference to the items they will be used to investigate and the questions they will be used to answer about them. I don't think this is true of functional imaging. As I'll illustrate by comparison to OCM imaging, FI reliability must be evaluated relative to the particular question an image is used to answer in a way which sets functional images apart from the kinds of evidence Woodward's account applies to most straightforwardly.

ii. Functional images are brain maps, colored to indicate BI levels estimated for various anatomical regions—reds for high levels, purples for low levels, intermediate colors for intermediate levels. PET scanning typically employs a radioactive oxygen isotope, O^{15} , introduced into the bloodstream as a component of H_2O^{15} , to produce the data the BI are calculated from. Scanning begins once it has had time to reach the brain. Positrons emitted by decaying O^{15} travel a few millimeters, encounter, and interact with electrons to produce paired gamma rays which travel out of the brain in opposite directions along a straight line toward the scanner.² In fMRI, manipulations of a magnetic field induce oxygenated and deoxygenated hemoglobin to emit weak radio waves of measurably different properties. In both cases, the estimation of BI levels and locations depends upon the fact that blood volume increases in the neighborhood of increased neuronal activity. Since there is more O^{15} where there is more blood, more gamma rays should be produced near regions of elevated neuronal activity than near other brain regions. With regard to fMRI, because more oxygenated hemoglobin circulates to neuronally active structures than their activity consumes, the level of oxygenated hemoglobin is elevated downstream from them.³ To produce a functional image, what I'll call a BI image must be computed from the radiation or radio signals the scanners have detected, and projected onto what I'll call an anatomical image. This is a sloppy way of saying that a mathematical representation of colors computed from the radiation or radio signals must be used to assign colors to pixels on a brain map. Here are a few of the features of this process which make FI evidence epistemically peculiar.

①Computed as it is from radiation or radio signals which fluctuate while anatomical structures remain the same, the BI image contains little anatomical information. However, the locations of the posterior and anterior commissures, along with the left, right, top, and bottom extremes of the brain are indicated clearly enough for use in projecting the BI onto an anatomical image. The BI

2 *For details see....

3 *for details see....

image is oriented on a coordinate system whose axes intersect one another at the anterior commissure. The y axis runs from front to back through the posterior commissure; the x, from left to right; and the z, from top to bottom. Numerical specifications of the positions of colors assigned to BI along these axes are used to assign them to regions of the anatomical image. The most commonly used anatomical image has been the Talairach–Tournoux atlas, a sequence of detailed maps derived from post mortem photographs of slices of a single brain. Although the BI image can be projected directly onto the Talairach–Tournoux atlas, many investigators project it first onto an MRI image⁴ of the brain from which it was derived. Others project it onto a template produced by computationally averaging MRI images of a number of different brains. The result is often warped onto Talairach–Tournoux atlas. Because the geography of the atlas differs observably from that of any individual brain, some investigators use estimates of the probability that a given location on the anatomical image corresponds to any given structure in an actual brain in reasoning from functional images.⁵

One reason commonly given for using a standard anatomical map like the Talairach–Tournoux atlas is that claims about functionally significant anatomical features cannot be satisfactorily generalized if they are too closely tied to the geographical idiosyncrasies of one, or a very few brains. Many important landmarks are fissures, furrows, or bulges of neuronal tissue which differ from brain to brain according to the way cortical tissue folds during embryological and later development. Furthermore, even if brains could be flattened to correct for these variations, observably large differences would remain.⁶ (Caviness Jr., et al.[1996], p.566) Another reason is that it is easier to compare BI patterns from different experiments if they are projected onto the same anatomical image. (Bavelier et al. [1997] p.671)

The Talairach–Tournoux atlas is remarkable for its unfaithfulness to the geographies of individual brains. Produced as it was from photographs of sections of the brain of a 60 year old woman, it departs in detail from male brains of all ages, and female brains of many ages. Furthermore, despite functionally significant anatomical differences between the hemispheres, photographs of the left brain were reversed and used to represent the right brain! (Collins [1998]

4 MRI images are computed from radio signals which vary with anatomical features instead of activity levels.

5. *To get a rough idea of what this involves, there are websites at which you can run a cursor over an atlas for one coronal(?) section of the brain, or select measurements on the Talairach coordinates, and see displays of the probability that your selection corresponds, e.g., to a given gyrus or subcortical structure.*(give addresses)

6 For example Watson et al.. identified visual area, V5 by '... determining the region of maximal rCBF change' in '... a region of the occipital lobe in which cortical activation by visual motion might be expected in each individual'. They found that the geographical location V5, so defined, differed observably from brain to brain. (Watson et al..[1993] p.83).

p.1) Furthermore, the geographical details of individual brains differ from those of an averaged anatomical image. These are two reasons to worry whether what appear as anatomically or functionally significant details in the image are really artifacts of the simplifications and idealizations which go into its production.

Functional image colors are computed pixel by pixel. PET imaging is so noisy that for a typical run, the standard deviation for distributions of

... differences between pixel values [over the threshold level] of two scans of the same subject performing the same sensorimotor control task... [was] 15–16%. (Haxby et al. [1991], p.542)

To correct for this BI images taken from different subjects (to save time and reduce the amount of radioactive material used for each subject) are statistically averaged. (Steinmetz and Seitz [1991] p.1158).⁷ Averaging washes out details of individual scans to bias PET images toward high BI levels, and anatomical structures which are relatively large, or which don't vary much from subject to subject. (*ibid*) Although inter-subject averaging is not required for fMRI noise reduction intra-subject averaging is required to correct for BI magnitudes which are idiosyncratic to one or just a very few performances of the same task, and inter-subject averaging is used to correct for the anatomical, psychological and physiological idiosyncrasies of individual subjects

Threshold levels BI are set to correct for the influence of baseline and other extraneous neuronal activity, of cognitively irrelevant blood circulation effects, etc. In the absence of precise estimates of the magnitudes of the relevant errors, investigators choose thresholds they hope will avoid significant mistakes without filtering out enough information to make the image useless.

Correlations between radiation and radio signals, BI activity, and CSNA are imperfectly understood—particularly with regard to their time courses. It is known that they begin and end at different times, and develop at different rates. It is not known precisely what the differences are. Furthermore, while it takes 40–60 seconds for PET, and 2–6 seconds for fMRI to collect enough signals to produce an image, CSNA supporting a cognitive function of interest may last only a fraction of a second. (Corbetta [1998] p.98, Haxby [1998] p.125) The color of a pixel must be adjusted to take account of all of this if it is to be indicative of CSNA of interest, rather than just the average of signals collected to produce the image. To this end, investigators choose a reference wave form to represent the development of CSNA beginning at the time of a sensory stimulus in response to which the subject initiates a mental task. Once pixel colors have been corrected for various artifacts (e.g., head motion during fMRI scanning) they are adjusted against the reference wave form. Sometimes a square waveform is used for mathematical convenience. Sometimes waveforms are adjusted to the results of human and animal experiments involving

⁷ In PET averaging groups of six to eight different subjects are considered small. (Watson, et al. [1993] p.;91

detections of neuronal activity associated with optical cortex and motor responses to visual stimuli. The adjusted wave forms can be expected to be more realistic than the square wave form, but the adjustments are only partially constrained by what is known about the relevant cognitive and neurophysiological phenomena. (Clare [2000] § 6.1, 6.3.2, Haxby [1998] p.129)

We've seen that the anatomical image typically represents the geography of an ideal, not a real brain. Threshold setting, image averaging, and choices of reference wave forms are among the further reasons why the levels and locations of BI functional images portray can be expected to differ from the levels and locations of BI and CSNA of any real brain.⁸ But for the most part, there are no reliable quantitative methods for calculating the extent of the differences.

iii. What's epistemically good about functional images is not that they are highly accurate with regard to BI levels or locations of individual brains. They aren't. What's good about functional images is not, as Mayo's account would suggest, that the net effect of the statistical operations which go into their production is to bring error-ridden, inaccurate BI estimates closer to what would have resulted from ideal experiments shielded from significant sources of error. Image averaging and the use of standardized anatomical maps are just two of the operations which take imaged BI levels and locations farther from, rather than closer to, those of any real brain. Furthermore, because it's unclear just what level the BI must reach to indicate CSNA, investigators cannot be confident they have avoided mistaking baseline for significant activity levels unless they set pixel thresholds high enough to risk filtering out some cognitively significant BI. And although it's plausible that adjusted reference wave forms introduce less error due to disparities between radiation and radio signal, BI, and CSNA time courses than square wave functions, their error characteristics cannot be precisely determined. Thus even though errors in signal detection are corrected for, and BI levels near baseline are not represented as cognitively significant, a number of computations involved in the production of a functional image increase certain kinds of error and take the imaged BI levels and locations farther from, rather than closer to those of individual brains.

The application of Woodward's counterfactual dependency characterization of reliability to functional images is complicated by the lack of techniques for calculating and reducing the inaccuracies introduced at various steps in their production. We can appreciate this by comparing FI to OCM imaging which allows (as FI does not) for estimations and improvements of accuracy.⁹ OCM uses a Michelson interferometer arranged so that light scattered back from structures on a focal plane (as much as 2mm deep) along a sample path encounters a reference beam whose length is adjusted by moving a mirror as required to produce interference fringes.

⁸ For more reasons, see Bogen [forthcoming].

⁹ The following sketch is based on an NSF application for OCM development functioning (1996) and a progress report (1997) by Richard Haskell et al. I am greatly indebted to Haskell for providing this material and answering questions about it.

Because the fringes are produced by interference which occurs only when the sample and reference paths are nearly the same length, the equipment can be adjusted so that only light scattered back from structures of interest can have a measurable effect on them. OCM images are computed from interference fringe measurements. Like FI image computation, this involves biological, physical and mathematical assumptions. But the computational decisions which embody them are far more heavily constrained by well established theory than those of FI imaging. And unlike FI, much of what is left open by theory can be checked in other ways. For example, Monte Carlo simulations were used to determine how long the sample beam must remain focussed on a given point to provide enough photons for a well resolved image, to figure out how to correct for the effects of multiple scattering, and to estimate differences in the way light would be scattered e.g., by organelles and by interfaces between intracellular and extracellular surfaces. To check the results, investigators can use them to predict images of artificial objects produced to ensure that their sizes, shapes, and optical features are well understood, and then make real images of them to compare the predictions to. Reliability can also be assessed by comparing OCM images to ordinary light photomicrographs of dead tissue sliced thin enough for adequate resolution. As these examples illustrate, OCM production can be calibrated and checked, detail by detail, against well established theoretical principles and clearly ascertainable features of natural and artificial structures. In keeping with Woodward's account of reliability assessment, investigation of the details of OCM image production can deliver a great deal of information about what could reasonably be expected of images produced under various counterfactual conditions.

None of this is possible for functional imaging. The relations between signals, BI, and CSNA are not well enough understood to enable investigators to make any such fine grained, quantitative determinations of the error characteristics of functional images by checking assumptions involved in their production against well established theory. Investigators cannot produce artificial BI or CSNA whose levels are well enough known to use to calibrate functional images. Nor can they independently measure naturally occurring brain activity as required for calibration.

How then can FI reliability be evaluated? Since it cannot be checked by appeal to theory, or independently known BI levels or locations, I suggest we think of functional images as accurate representations of the anatomy, and the cognitively significant BI of, imaginary, ideal, working brains. Just as the epistemic value of an OCM image depends not only upon how well it represents the relevant features of the cells whose image it is, but also upon how representative those cells are, I think the epistemic value of a functional image depends upon whether the ideal brain it portrays is representative of real brains with regard to the anatomical, physiological, and psychological factors relevant to the question under investigation. If this is correct, the epistemic evaluation of a functional image is a matter of determining what sorts and degrees of

resemblances between real brains and the ideal brain it portrays are required to answer the question under investigation, and finding out whether it is safe to assume those resemblances obtain.

For example, in 1998 Rugg used fMRI results obtained by Brewer and Wagner to support the hypothesis that (h) an event is more likely to be remembered if the subject attends to '... its meaning rather than to more superficial attributes (such as physical appearance)'. Rugg [1998] p.1151) Brewer, and Wagner made fMRI images of subjects performing sorting tasks involving pictorial or verbal displays they were not told they would have to remember. Later, they were given recognition tests were. For some pictures, functional images made as they were displayed during the sorting tasks indicated above baseline levels of neuronal activity in the right parahippocampal cortex and a small region of the right prefrontal cortex. Subjects tended to remember they had seen those pictures, and not to confuse them with items they had not seen earlier. Subjects made significantly more mistakes with pictures whose display during sorting appeared in the functional images not to have occasioned elevated activity in the parahippocampal and prefrontal regions. (Rugg [1988] p.1151, Brewer [1998] p.1186) Similar results were obtained for the left parahippocampal and prefrontal areas for verbal displays. (Rugg [1988] p.1151, Wagner [1998] p. 1189) In another experiment, when subjects sorted words and pictures in a number of different ways, functional images indicated higher levels of CSNA in these same regions of the brain when the sorting task involved semantic meaning than when only physical features were involved.(Rugg [1998] p.1151, Wagner [1988] pp.1189–90)

To produce his evidence, Brewer averaged BI images from 6 subjects and assigned colors to an averaged anatomical map to indicate BI levels above the thresholds he had chosen. He then 'normalized' the results 'to the stereotactic space of Talairach and Tournoux.' (Brewer et al [1998] n.13, p.1187). I suppose Brewer's images reliably indicate CSNA in precisely defined regions of an imaginary, ideal brain, and that the epistemic value of his evidence depends upon what (if anything) the activity in its hippocampal and prefrontal regions corresponds to in the brains of his subjects. The same holds for Wagner's evidence. Thus the value of the images depends [a] upon whether the parahippocampal and prefrontal regions of the ideal brain they portray correspond to functionally significant anatomical units of real brains, and if so, [b] whether the BI levels of the ideal brain are reliable indicators of CSNA in the corresponding areas of real brains. As said, averaged images projected on to the Talairach–Tournoux atlas cannot be expected to accurately represent BI or CSNA levels or locations of real brains. And in contrast to OCM, there is no well established theory of detailed dependencies between signal detections, BI levels, and CSNA by which to assess reliability with regard to [a] or [b]. Nor can reliability be evaluated by comparing levels and anatomical locations of BI displayed in the functional image to independently established levels or locations in real brains. Nor can reliability be established by calibration. Lacking the means, available for OCM, to evaluate and improve reliability by

investigating and fine tuning the imaging process, here is what can be done.

With regard to [a], consider the parahippocampal cortex. Postmortem examinations of humans, and dissections of primates establish that this is an anatomically discrete part of the brain, located in the medial region of the temporal lobe. With regard to its functional significance, clinical observations and neuro-cognitive psychological experiments have established that

... bilateral damage to the medial temporal lobe ... yields... a pervasive memory deficit for all new events and facts,

that left side lesions impair verbal memory, and right side lesions, nonverbal memory. (Brewer et al [1998] p.1185) The damage is typically too extensive for investigators to establish precisely the locations of the structures whose function is required for memory—let alone where those structures are relative to a Talairach–Tournoux or averaged brain image. But the lesion evidence does argue that in real brains the medio-temporal regions which contain the parahippocampal cortex are functionally significant in the sense that memory requires their normal functioning. Animal experiments provide further support. Parahippocampal lesions in monkeys cause poor performance in memory tasks analogous to those which humans with medio-temporal lesions perform badly.(Bauer et al [1993] p.552). Thus there is reason to think real brains resemble the ideal brain of Brewer imaged in that that their medial temporal lobes contain anatomical structures which support memory. There is no doubt that real and idealized anatomical structures differ from one another with regard to their shapes, relative sizes. Even though the differences cannot be precisely determined, it is still plausible that however they may differ in other respects, the parahippocampal cortices of real and ideal brains are, or are located in, or are near, anatomical structures whose activity contributes to memory. Similar arguments apply to the prefrontal region.

With regard to [b] if the parahippocampal and prefrontal areas of individual brains (or the regions which contain them) are functionally significant with regard to memory, does their activity resemble that of their idealized counterparts as needed to argue for **h**? Rough, qualitative evaluations of the images relative to the question Brewer, Rugg, and Wagner were interested in are enough to make it plausible that the required resemblance holds. Their question was whether the parts of the brain whose normal functioning is known (from the clinical and animal experimental evidence just mentioned) to be necessary for memory are the same as those whose activity above baseline levels correlated with good performance on the recognition tests, and which were active above baseline levels when subjects attended to the meanings of displayed words and figures. Neither precise quantitative measures of their activity, nor accurate specifications of sizes, shapes, or relative positions of the relevant brain regions, nor accurate, quantitative estimates of functional imaging accuracy are required to justify an affirmative answer to this question. Even though image averaging and projection onto the Talairach–Tournoux atlas

obscure the geography of the brain, there is enough independently obtained anatomical evidence to show that individual brains contain medio-temporal and prefrontal regions which correspond to the parahippocampal cortex and the relevant prefrontal area of the ideal functional image brain. There is no way to determine quantitatively the inaccuracies and imprecisions produced by the mathematical manipulations and assumptions involved in functional image production. But thresholds can be set high enough to warrant confidence that above baseline BI and CSNA levels in the parahippocampal and prefrontal cortices of the ideal brain correspond to above baseline levels of BI and CSNA in the corresponding areas (whatever they may be) of the imaged subjects' brains. In setting them high enough for this purpose, the investigator sacrifices accuracy to insure that cognitively insignificant BI levels will not be imaged as indicative of significantly elevated CSNA.

This example suggests that the evidential value of a functional image depends upon [a] whether regions of interest in the ideal FI brain correspond, as required for the investigation of the question of interest, to functionally significant regions of real brains, [b] what particular sorts of inaccuracies and imprecisions in image production and computation of BI levels and locations would make it impossible in practice to reason cogently from the image to an acceptable answer the question of interest, and [c] whether the image was produced in such a way as to rule out those mistakes without introducing others which would make it epistemically unsuitable for the investigation of that question. The evaluation of a functional image along these lines may involve Mayonian uses of error analysis, and the establishment of Woodwardian counterfactual dependencies. But at least as things stand now, an analysis of image production will not deliver quantitative determinations of error characteristics, or degrees of imprecision and inaccuracy like those available for the evaluation of OCM and many other sorts of experimental evidence.

Bibliography

- Bavelier et al. [1997]—D Bavelier et al., 'Sentence Reading: A Functional MRI Study at 4 Tesla', *Journal of Cognitive Neuroscience*, 9:5
- Bogen [forthcoming]—J Bogen, 'Functional Imaging—Some Epistemic Hot Spots' in R Grush, et al, *Limits of Science: Philosophical Problems in the Neurosciences*, Pittsburgh: University of Pittsburgh Press
- Bauer et al [1993]— RM Bauer, B Tobias, E Valenstein, 'Amnesic Disorders' in K Heilman, E Valenstein, eds., *Clinical Neuropsychology*, New York: Oxford University Press.
- Caviness Jr, et al.[1996]— VS. Caviness Jr. et al., 'MRI-Based Topographic Parcellation of Human Neocortex: An Anatomically Specified Method with Estimate of Reliability', *J Cognitive Neuroscience* 8:6
- Clare [2000]—S Clare, *Functional MRI Methods and Application*, PhD thesis
<http://www.fmrib.ox.ac.uk/~Stuart>
- Collins [1998]—J Collins, 'Probabilistic Atlases of Cerebral Anatomy',

<http://w Hlt491752737w Hlt491752737w- Hlt491752476 Hlt491752476 sop.inria.fr/epidaure/pe... dain/neuro98/abstracts/collins.html>.

(Corbetta [1998]—M Corbetta, 'Functional Anatomy of Visual Attention in the Human Brain' in Parasuraman [1998]

Haxby [1998] p.125 JV Haxby et al, 'Functional Magnetic Resonance Imaging and the Study of Attention' in R Parasuraman [1998]

Mayo [1996]—D Mayo, *Error and the Growth of Experimental Knowledge*, Chicago: University of Chicago Press.

Parasuraman [1998]—R Parasuraman, ed., *The Attentive Brain*, Cambridge: MIT Press

<<5468>>