科思传播
生命科学
·导读版·

STATISTICAL PARAMETRIC MAPPING: The Analysis of Functional Brain Images

统计参数图：脑功能成像分析

Karl Friston, John Ashburner, Stefan Kiebel,
Thomas Nichols and William Penny

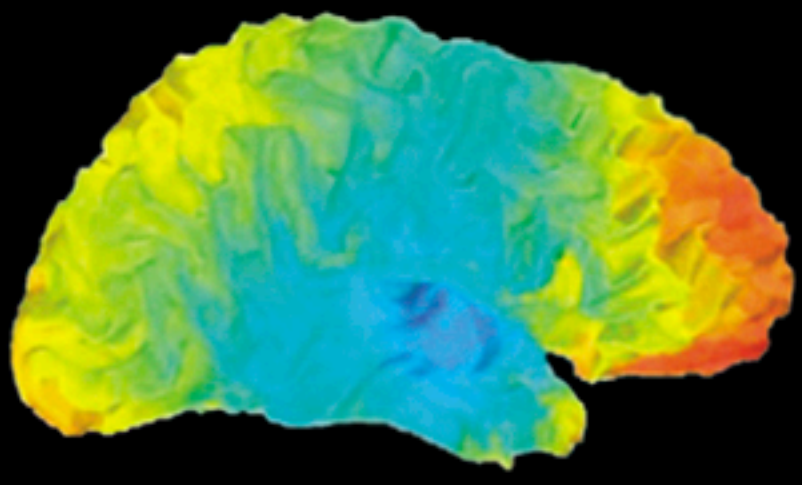

ELSEVIER
原版引进

科学出版社
www.sciencep.com

STATISTICAL PARAMETRIC MAPPING:

The Analysis of Functional Brain Images

统计参数图:脑功能成像分析

Edited by
Karl Friston
John Ashburner
Stefan Kiebel
Thomas Nichols
William Penny

科学出版社
北京

图字:01-2009-7227 号

This is an annotated version of
STATISTICAL PARAMETRIC MAPPING: The Analysis of Functional Brain Images
Edited by Karl Friston, John Ashburner, Stefan Kiebel, Thomas Nichols and William Penny.

ISBN 13: 978-0-12-372560-8

Authorized English language reprint edition published by the Proprietor.
ISBN 13: 978-9-81-272558-5

Elsevier (Singapore) Pte Ltd.
3 Killiney Road
#08-01 Winsland House 1
Singapore 239519
Tel: (65) 6349-0200
Fax: (65) 6733-1817

First Published 2010
<2010> 年初版

图书在版编目(CIP)数据

统计参数图:脑功能成像分析=Statistical Parametric Mapping:The Analysis of Functional Brain Images:英文/(英)弗理斯顿著.—北京:科学出版社,2010

ISBN 978-7-03-026349-0

Ⅰ.统… Ⅱ.弗… Ⅲ.脑-功能-成像-英文 Ⅳ.R338.2

中国版本图书馆 CIP 数据核字(2010)第 003496 号

责任编辑:田慎鹏 李小汀/责任印制:钱玉芬/封面设计:耕者设计工作室

科学出版社 出版
北京东黄城根北街 16 号
邮政编码:100717
http://www.sciencep.com

双青印刷厂印刷
科学出版社发行 各地新华书店经销

*

2010 年 1 月第 一 版 开本:787×1092 1/16
2010 年 1 月第一次印刷 印张:41 3/4 插页 16
印数:1—1 000 字数:951 000

定价:168.00 元

(如有印装质量问题,我社负责调换〈双青〉)

导　读

一、脑功能成像概观

“科学探索的前沿由两个因素决定：一是观测和实验的工具，一是概念上的革新；两者缺一不可。”[①]这条真理适用于几乎一切科学领域，特别是像认知神经科学这样新兴的前沿学科。近十年来，认知神经科学的蓬勃发展，不可否认地得益于心理学、认知科学、神经科学和信息科学等领域相互交叉融汇导致的观念创新以及实验手段的精准化和多样化的巨大推动。其中，脑成像技术的成熟和普及功不可没。

“很少有哪些科学进展可以像为正在工作的人类大脑进行‘拍照’这样引人注目”[②]。人们为什么如此着迷于脑功能成像？因为，为了理解正常人类大脑的工作机制，功能成像是必不可少的。它试图将不同的心理过程定位于大脑的不同部分，也即绘制一张标明哪些区域负责哪些加工过程的“脑功能解剖图”。如同“人类基因组计划”一样，对人类大脑功能的探索寄托着人们“认识自我”的亘古不变的理想和愿望。

在众多功能成像技术中，功能性磁共振成像(functional magnetic resonance imaging，fMRI)以其技术上的优越性和相对较低的成本脱颖而出，成为了应用最广、成就最显著的一种。fMRI 在技术上的优越性体现在以下几个方面：

1. 完全的无创和非侵入性。在脑功能成像技术成熟之前，认知神经科学的发展受到伦理上的很大限制：研究者不能像在非人类动物身上那样使用穿透性的探测设备，如微电极，探测正常人类受试者的神经活动。这就使得研究者不能确定，人类认知功能(如视觉加工、记忆、情绪加工等)的神经机制是否可以由在动物上得到的神经生理模型加以解释。另一方面，更受到关注的人类特有的高级认知功能，如语言、社会交往等，是无法在实验动物身上得到确定性的信息的。早期的功能成像技术，如正电子断层扫描(Positron Emission Tomography，PET)，向着解决这一矛盾迈出了探索性的第一步，并取得了一定的成就。然而，由于 PET 技术需要为受试者注射放射性示踪剂，因此受试者在一次实验中接受的实验时间受到了很大限制，这直接制约了 PET 技术所能研究的问题的范围和所提供证据的准确性。而 fMRI 作为一种完全无创和非侵入性的成像技术，不仅解决了“伦理困境”，而且极大地丰富了功能成像的研究对象。由于 fMRI 采集的是神经生理活动自主产生的信号，实验时程理论上不受到技术和安全性的限制，同类刺激可以多次呈现给受试者，这就将数据的准确性和可靠性提升到一个新的高度。

2. 较高的时间和空间分辨率。由于人们关心的认知和神经加工过程是非常迅速的，通常在秒甚至亚秒量级；同时，大脑中功能相近的神经元所聚集成的模块的大小一般也不到毫米量级。这就要求探索认知神经机制的仪器有相当的时空分辨率。PET 技术的信

---

① Gazzniga, Ivery and Mangun. *Cognitive Neuroscience: The Biology of the Mind (3rd edition)*, Chapter 4.

② Huettuel, Song and McCathy. *Functional Magnetic Resonance Imaging*, Chapter 1.

号依赖于放射性示踪剂的半衰期,一般在分钟水平,而空间分辨率在厘米量级;事件相关电位(ERP)技术的信号依赖于神经元集群放电,时间分辨率可以到百毫秒量级,但空间分辨率较低。而目前常用的高场强(1.5～3 特斯拉)MRI 扫描仪的时间分辨率可达到秒量级,空间分辨率达到毫米水平。

3. *成熟而多样的数据分析手段*。经过十多年的理论探索和实践应用,认知神经科学和相关领域的研究者已经发展出一套成体系的数据分析手段,适合于回答各种类型的问题。最基本的单变量统计检验(通用线性模型)和与之相应的感兴趣区域分析,虽然已经足够好的给出不同认知活动间神经活动的差异,但它们已不能满足研究者日益深入的探索。各种复杂精细的数据处理方法应运而生:可以捕捉多神经元群体活动模式的"多变量模式分析(Multivariate Patten Analysis, MVPA)",探索不同脑区在功能上的相互联系的各种功能连接模型(如 Dynamic Causality Model, Granger Causality Model)等等,已经逐渐成为认知神经科学家的"常规武器",为理解人类大脑功能提供了越来越丰富和深刻的证据。

4. *与其他神经科学手段的融合*。尽管 fMRI 信号来源的神经生理学根据还基于一些尚未得到肯定的假设,但 fMRI 信号本身与其他神经科学手段,如单细胞和多细胞记录、局域场电势,以及事件相关电位和事件相关电位场(ERF)得到的证据非常好的吻合。而近年来一个引人注目的发展趋势即是 fMRI 技术与其他手段的融合,如 fMRI 与 ERP 同时记录,fMRI 数据与行为数据的相关,fMRI 反应模式与受试者人格倾向、生理化学指标和基因类型的相关等。这些证据的相互印证不仅使它们更加确定,更重要的是,技术的融合使研究者发现了单独使用任何一种手段所不曾发现的现象。这极大地推动了认知神经科学领域的发展,深化了人们对脑功能的认识。

无论在科学探索的哪个领域,一项好的研究都至少包括以下几个要素:一个有创新型的想法或问题,一套缜密合理的实验设计,一组合适的实验仪器或手段,以及一系列严格而巧妙的数据处理方法。在脑功能成像领域,"统计参数图(Statistical Parametric Mapping, SPM)"这一成体系的数据处理方法当之无愧的是第四方面的杰出代表。该系统以数学和生理物理模型为基础,为为数众多的认知神经科学研究者(他们中的大多数也许并不精通数学和生理学)搭建了从成像数据到神经活动之间的桥梁。从功能成像技术诞生之日起,SPM 就追随并推动着这项技术的发展和成熟。直到今天,当我们翻开功能成像研究的期刊,就会发现 SPM 仍是功能成像数据处理的主流手段。

近年来,我国认知神经科学也开始蓬勃发展。本书在中国的引进出版,将进一步推动中国的认知神经科学,特别是功能成像研究的发展。

二、本书各章简介

本书共七部分,第一部分是导论,统领全书的内容;第二到第四部分讲述脑成像数据分析的关键步骤:图像变换、建模和推断;第五到第七部分讲述神经元响应的生物物理模型,这些部分与功能连接和有效连接的分析紧密联系在一起。

第一部分　导论　对后面要讲的内容作简要的概括,使读者把握全书的精要。首先在第一章中简述了 SPM 的简史。在第二章中,总结了进行脑成像数据分析的基本方法和程序,综述了数据分析的三个步骤——图像变换、建模和推断——中主要的问题,包括:

(1)处理排除被试的头部运动对数据带来的干扰,方法是利用刚体变换和非线性配准将这种干扰最小化;(2)空间的标准化,把每个被试的数据标准化到同一空间,便于被试间的比较;(3)数据的光滑化;(4)如何应用通用线性模型;(5)利用随机场理论进行统计上的拓扑推断;(6)实验设计的相关问题。第三章总结了在最后三部分中讲述的模型,包括解剖模型、统计模型、生物物理模型。

第二部分　计算解剖学　本部分讲述对图像进行预处理的方法,包括刚体配准、非线性配准、图像分割和形态计量学。

刚体配准:图像配准(alignment)在功能成像分析的许多方面都有重要意义,刚体配准是图像配准中最简单的一种。在脑成像研究中,血液动力学变化产生的信号变化会被被试者头动造成的信号干扰。有时,这种干扰信号甚至远远强于真正感兴趣的信号。在实际试验过程中,即使尽量避免被试头动,被试在扫描时的头动依然不能完全避免,因此在预处理时需要进行头动的校正。大量扫描会造成微小的系统误差的积累,变得非常显著,从而误差项扩大,对数据分析造成干扰。若没有进行合适的校正,由试验范式引起的被试活动可能会人为的造成激活的假象,从而会误导研究者做出错误的推断。

运动校正的重要性还在于可以增加统计敏感度。统计检验一般是以 t 检验为基础检查信号的改变,t 值的计算包含了信号差异造成的残差,如果活动的假象被包含在残差中,那么就会使 t 值变小,从而降低实际激活检验的敏感度。

刚体配准还可应用于形态计量学领域,将不同时间获得的单个受试者的图像进行比较,从而判断其图像模式的变化。造成这种变化的原因可能有很多种,但是经常与病理机制有关。因为扫描的都是同一个被试,所以这种分析的第一步就是利用刚体配准将图像一起配准。

图像配准的最简单用途是估计一对图像的映射。方法是:确定参考图,其他的图像(浮动图)空间变换后与参考图匹配。变形浮动图匹配参考图,主要是确定参考图中每个体素位置匹配到浮动图中相应位置的映射,浮动图在新位置进行二次重采样,而映射可以看成估计变换参数的一组参数。三维空间的刚体变换一般由六个参数决定:三个平移参数和三个旋转参数。图像配准通常分两步:第一步估计配准参数,第二步利用这些参数把浮动图移动到参考图的坐标空间。通常情况下,配准参数的估计是通过迭代方式完成的。

非线性配准:很多时候仅仅用刚体配准是不够的,因为被扫描对象的形状可能发生变化,这种变化可能由大脑的发育、衰老或者疾病造成。非线性配准在 SPM 中的最主要应用是空间的标准化。很多时候,人们希望把每个个体的图像大致的变形到脑的标准空间,即所谓的空间标准化,这样方便在不同的被试之间进行信号的平均,从更广泛的意义说,方便了信号的整体处理,以便提取出更多的信息。功能成像中,其作用主要有两个:首先是对确定在被试中普遍发生的现象有重要作用,另外通过标准空间的欧几里得坐标可以方便地报告激活发生的位置。

图像配准的方法大体有两类:基于标记配准和基于体素灰度值配准。基于标记配准可以先判断出参考图和浮动图之间的相似特征(即标记),再找出最优变形的方法将两者重叠。这些标记可以是点,也可以是线和面。基于体素灰度值配准可以优化浮动图和参考图之间的体素相似性,鉴别空间变形方法,此时图像被看成是没有标记的连续物。匹配的标准通常是最小差异平方和或最大图像相关。近来出现了新的综合方法,把以上两类

方法结合起来使用。

图像分割:图像分割的目的是将 MR 结构像分割为不同的脑组织。这些脑组织是灰质、白质和脑脊液。分割的方法基于改进的高斯混合模型和脑组织空间分布的先验概率知识。先验概率可通过优化方法和实验数据所提供的信息被进一步修正。正常的脑结构 MR 图像按组织划分为三部分:灰质、白质和脑脊液,可手动将高质量的 T1 图像进行划分,操作方法是选择合适的图像强度范围值,此范围值要包括某组织中大多数体素的灰度值。图像分割的方法通常有两种:或者用组织分类的方法,或者用图像变形的方法,把一个作为模板的大脑变形以和要分割的大脑相配。本章介绍了一种将两种方法结合在一起的框架。

形态计量学:形态计量学主要用统计推断来描述多个被试间脑结构差异,或寻找脑形状的相关信息以说明疾病的严重性。形态计量学方法有多种,VBM 是被广泛应用的一种。跟其他方法不同的是,VBM 并不直接涉及到解剖;它将图像视为连续的标量度量,检验在适当的空间尺度下的局部差异。这个尺度是受光滑性控制的,光滑性在 VBM 中是至关重要的,因为它为选择适当的尺度提供了下界。VBM 可以有效的在尺度空间中寻找解剖差异。VBM 广泛流行的原因有数个:首先,与传统解剖结构的分析不同的是,它可以检验大脑任何位置的差异;其次,它可以用于大脑解剖学的任何度量方式。恰当的选择数据和数据变换,大量的解剖特征可以通过简单无偏的方法加以分析。

第三部分　通用线性模型　本部分讲述通用线性模型的主要理论。通用线性模型是分析功能成像数据的重要工具,经典的脑功能数据分析几乎都是基于这个模型,主要包括三个方面:模型建立,参数估计,统计推断。

通用线性模型假定了实验和观测数据之间的特定联系。这种联系包含在特定的设计矩阵(design matrix)中,为了检测某种特定的联系,需要建立统计上的对照(contrast)。对同一设计矩阵应用不同的对照,可以检测多种效果而不必更改模型。

利用通用线形模型进行统计推断时,经常要利用 F 检验,该检验的有效性依赖于协方差结构的匀质性,即所谓的“球形假设”,但是当这个假设不成立时,就应该注意此时 F 检验的可靠性,若偏离严重,则应该采用新的方法来解决这个问题,比如设限的最小二乘法。

随机效应分析(Random Effect Analysis)关注如何对多个被试的脑功能成像数据进行推断。早期大量的脑成像研究将多被试的数据看成固定效应(Fixed Effect),这样最后做出的统计推断也是关于这批被试的,只能视为个例研究,不能做出对整个人群的推断。而随机效应分析从新的观点出发,将每个被试视为从人群中随机抽取的,这样被试的效应可以视为人群中满足某种分布的随机变量,从而可以做出对人群总体的统计推断。

通用线性模型和随机效应模型都可以视为分层模型的一种特例,可以纳入到更一般的分层模型中统一处理。分层模性若只有一层,就是通用线性模型;随机效应模型则是一种两层的分层模型。对分层模性的参数估计和统计推断可以利用参数经验贝叶斯法则,上一层的信息作为下一层的先验知识,进行参数估计。

处理 fMRI 数据需要利用一般的线性时不变系统(Linear time-invariant system)的理论,该理论是对通用线性模型的推广。在具体的建模过程中,BOLD 信号利用血液动力学函数的卷积来处理。卷积模型和时间基函数是处理 fMRI 数据的重要工具。BOLD 信

号的特殊性质使得实验设计成为一个复杂的课题,从信号处理、统计效度和回归因子的相关性等方面考虑可以给最大化实验设计的效度带来帮助。

第四部分　经典统计推断　本部分讨论模型参数的经典统计推断和检验方法,主要讲述了随机场理论及其在功能成像数据中的应用,以及非参数检验。

单变量神经图像的统计模型建立与参数估计给予在每个象素点的统计分析,所有的象素的统计量形成一幅统计参数图。研究者希望找出激活脑区,并希望第一类错误得到合理的控制。因此,必须在整个脑图上考虑统计检验,不能分别只对单个象素进行统计检验,这样就面临多重比较(Multiple Comparison)统计检验。由于脑成像数据中相邻点之间的信号相关性较强,所以传统的 Bonferroni 矫正方法要求过于严格,得出的统计结果不准确。而随机场理论(Random Field Theory)考虑了脑成像数据的空间相关性,因此利用其进行检验更准确。利用随机场理论可以进行两种脑区激活的检验:一是激活量的大小,二是激活量的体积。

在数据处理过程中,通常不会预先知道脑的哪部分会激活,所以需要对整个脑或其中的某一部分进行搜索,这给多重比较检验提出了问题。解决办法之一基于 T 检验或 F 检验,通过调整阈值选择统计值大的区域,选择阈值时可用随机场理论;第二种方法基于考察统计值超过阈值区域的空间范围,这种方法的理论基础也是随机场理论。

此外,非参数统计检验也已经进入脑成像的研究中,虽然还不是主流。参数检验依赖于分布的假设,大多数情况下都有正态性假设。当这种分布假设偏离严重的时候,统计检验就会出问题。非参数统计的优点在于不需要做分布的假设。一种称为非参数排列检验的方法已被应用到不少脑成像的研究中。在现在的功能成像实验中,由于需要进行很多的假设和近似,所以非参数统计的方法经常表现得要比参数方法好些。

第五部分　贝叶斯推断　本部分讲述了贝叶斯方法进行功能成像数据的统计推断。强调了贝叶斯方法和经典方法的共同点。实际上,经典方法可以被纳入到经验贝叶斯推断的框架当中,两者在算法上的连接点是分层线性模型。经验贝叶斯方法用上一层的参数估计结果作为下一层的经验的先验概率分布,因此它与传统贝叶斯推断是有区别的,传统的贝叶斯方法会事先选择恰当的先验函数作为先验分布。经验贝叶斯方法涉及的参数主要是方差和协方差,它们的估计都是通过 EM 算法实现的。

目前的脑成像领域的研究,脑成像的推断大多被限制在 SPM 的传统推断方法,但是经典推断方法是有缺陷的,在给定充足的数据或者敏感性的时候,与零假设的平凡的偏离也会被认为是显著的。经典推断方法之外的一种选择是在给定数据的情况下,利用后验概率分布进行推断,这就是所谓的 PPM 方法。PPM 不会遇到经典推断中难以解决的多重比较统计检验问题,因此可能发展为脑成像领域一种相对更有力的推断方法。利用经验贝叶斯理论可以方便的建立后验概率分布图。在具体的计算实现过程中,可能会遇到一些困难,可以使用变分技术来满足计算上的可行性,EM 算法实际是这种技术的一个特例。利用通用线性模型和变分技术我们可以建立 fMRI 和 EEG 的时空模型,来刻画这两种过程的时空性质。

第六部分　生物物理模型　本部分主要关心的是神经响应的生物物理模型和怎样对这些模型的参数做出推断。与先前的内容相比,这一部分更关心观测到的功能成像信号是如何产生的,以及背后的物理和生理机制。首先从血液动力学讲起,然后论述 EEG 和

MEG 信号产生的电磁机制。后续章节考虑更深层的生理机制,通过神经元群体的平均场模型发展出事件相关电位的神经群体模型。最后考虑了在模型参数估计和模型选择中会遇到的一些一般性的问题。

事件相关功能核磁共振出现以后,非线性在 fMRI 激活响应中的重要性日益受到重视。这种非线性主要应用于刺激间的交互作用,可以增强或抑制由先前的刺激诱发的响应。前面章节曾介绍了一种不是基于机制模型的 Volterra 序数列表达法,它采用非线性系统识别的方法刻画刺激间的相互作用。Buxton 阐述了一种直觉上合理的 fMRI 血液动力学转换模型。随后 Mandeville 提供了对脑血流和血体积之间关系的理论和约束条件。综合以上的基于系统识别理论和动态机制模型理论,求证了 Balloon 模型在处理时间序列信号中的非线性是足够的,其模型参数在生物学上也是可信的。结论是 Balloon 模型可以产生出与试验近似的 Volterra 核。fMRI 和 EEG 的正演化模型都可以建立起来,不同的是前者更关注其中的时间信息,而后者更着重处理空间信息。利用经验贝叶斯方法可以实现 EEG 和 MEG 信号源的重建。

以上这些考虑的都是怎样从模型出发估计信号的源,相反的问题在本部分中也得到了细致的处理,即如何对神经动力学产生信号这个过程进行建模。在本部分,fMRI 和 EEG 背后的神经元动力模型被建立起来,此外一种在神经元层次将 fMRI 和 EEG 联合在一起的模型被提出来,为联合两种手段提供了一种潜在的途径。

第七部分　大脑连接　本部分讲述把脑作为一个完整的系统进行研究的概念和方法。首先介绍脑功能的整合,从拥有层次连接模式的脑皮层如何处理信息的角度介绍功能整合的神经生物学基础,为脑功能整合提供了研究框架。实际上,经验贝叶斯理论是研究以层次模式连接的脑皮层之间的相互作用的基础。然后介绍了功能连接和有效连接的定义、数学原理和应用。这两种方法都描述了脑的不同部分之间的关系,不过描述的角度不同。这些方法可以研究脑的相关部分在正常生理活动下的协同工作原理,也可以表示脑疾病对于脑的工作方式的影响。接下来的两章从两个互补的角度表述了脑功能整合模型,即 Volterra 核和动态因果模型(Dynamic Causality Model)所使用的状态—空间表示法。最后几章讲述了动态因果模型。

脑的结构组织所体现的功能特征主要体现功能分区和功能整合。功能分区的形成主要是由于脑中功能相近的细胞和组织的物理位置倾向于集中在一个小的区域,这使得某个认知功能的能量消耗及传输时间能够最小化。某个认知任务中的功能整合主要通过此任务中各个功能分区之间的有效连接来实现。功能连接定义为“空间上分离的部位,在神经生理活动中的相互关系”。功能连接为功能整合提供了初级的分析方法。有效连接描述神经活动中,一个脑区对另一个脑区施加的影响。脑中的有效连接会发生变化,在不同的认知任务中具有选择性。

动态因果模型从功能整合的角度研究脑,衡量其在不同实验条件下变化的有利工具。该模型通过实验数据的时间序列推断三个主要的参数:外部刺激对被刺激脑区的直接影响;实验任务下脑区的有效连接;实验条件对这些有效连接的影响。其主要作用是推断脑区之间的有效连接以及这些连接如何随实验条件而变化。需要注意的是,该模型并没有发现脑区连接的功能,研究者需要自己建立在神经生物学上正确的模型。动态因果模型的功能能否被完全利用,取决于实验设计,多因素(factor)实验设计可以被用来探索有效

连接的变化。

附录　附录部分对正文部分中用到的数学理论提供了较为严格的描述，包括：线性模型和推断，动力系统，EM 算法，拉普拉斯近似下的变分贝叶斯方法，卡曼滤波和随机场理论。其中涉及到比较多的数学推导，对数学推导细节不感兴趣的读者可以有选择的略过。但是这些数学背景知识在正文中经常使用，仔细阅读可以使读者更细致地理解本书的理论。

周晓林

2009 年 12 月

目　　录

（周晓林 译）

Contents

Part 6 Biophysical models

Part 7 Connectivity

Appendices

Acknowledgements

This book was written on behalf of the SPM co-authors, who at the time of writing, include:

Jesper Andersson
John Ashburner
Nelson Trujillo-Barreto
Matthew Brett
Christian Büchel
Olivier David
Guillaume Flandin
Karl Friston
Darren Gitelman
Daniel Glaser
Volkmar Glauche
Lee Harrison
Rik Henson
Andrew Holmes
Stefan Kiebel
James Kilner
Járámie Mattout
Tom Nichols
Will Penny
Christophe Phillips
Jean-Baptise Poline
Klaas Stephan

We are also deeply indebted to many colleagues who have developed imaging methodology with us over the years. Though there are too many to name individually we would especially like to thank *Keith Worsley* and colleagues at McGill. We are also very grateful to the many imaging neuroscientists who have alpha-tested implementations of our methodology and the many researchers who regularly make expert contributions to the SPM email list. We would especially like to thank *Jenny Crinion, Alex Hammers, Bas Neggers, Uta Noppeney, Helmut Laufs, Torben Lund, Marta Garrido, Christian Gaser, Marcus Gray, Andrea Mechelli, James Rowe, Karsten Specht, Marco Wilke, Alle Meije Wink and Eric Zarahn.*

Finally, any endeavour of this kind is not possible without the patience and support of our families and loved ones, to whom we dedicate this book.

Part 1

Introduction

CHAPTER

1

A short history of SPM

K. Friston

INTRODUCTION

For a young person entering imaging neuroscience it must seem that the field is very large and complicated, with numerous approaches to experimental design and analysis. This impression is probably compounded by the abundance of TLAs (three-letter-acronyms) and obscure terminology. In fact, most of the principles behind design and analysis are quite simple and had to be established in a relatively short period of time at the inception of brain mapping. This chapter presents an anecdotal perspective on this period. It serves to explain why some ideas, like *t*-maps or, more technically, statistical parametric maps, were introduced and why other issues, like global normalization, were crucial, even if they are not so important nowadays.

The history of human brain mapping is probably shorter than many people might think. Activation studies depend on imaging changes in brain state within the same scanning session. This was made possible using short-half-life radiotracers and positron emission tomography (PET). These techniques became available in the eighties (e.g. Herscovitch *et al*., 1983) and the first activation maps appeared soon after (e.g. Lauter *et al*., 1985; Fox *et al*., 1986). Up until this time, regional differences among brain scans had been characterized using hand-drawn regions of interest (ROI), reducing hundreds of thousands of voxels to a handful of ROI measurements, with a somewhat imprecise anatomical validity. The idea of making voxel-specific statistical inferences, through the use of statistical parametric maps, emerged in response to the clear need to make inferences about brain responses without knowing where those responses were going to be expressed. The first *t*-map was used to establish functional specialization for colour processing in 1989 (Lueck *et al*., 1989). The underlying methodology was described in a paper entitled: 'The relationship between global and local changes in PET scans' (Friston *et al*., 1990). This may seem an odd title to introduce statistical parametric mapping (SPM) but it belies a key motivation behind the approach.

Statistical maps versus regions of interest

Until that time, images were usually analysed with analysis of variance (ANOVA) using ROI averages. This approach had become established in the analysis of autoradiographic data in basic neuroscience and metabolic scans in human subjects. Critically, each region was treated as a level of a factor. This meant that the regional specificity of a particular treatment was encoded in the region by treatment interaction. In other words, a main effect of treatment *per se* was not sufficient to infer a regionally specific response. This is because some treatments induced a global effect that was expressed in all the ROIs. Global effects were, therefore, one of the first major conceptual issues in the development of SPM. The approach taken was to treat global activity as a confound in a separate analysis of covariance (ANCOVA) *at each voxel*, thereby endowing inference with a regional specificity that could not be explained by global changes. The resulting SPMs were like X-rays of region-specific changes and, like X-rays, are still reported in maximum-intensity projection format (known colloquially as glass-brains). The issue of regional versus global changes and the validity of global estimators were debated for several years, with many publications in the specialist literature. Interestingly, it is a theme that enjoyed a reprise with the advent of functional magnetic resonance imaging (fMRI) (e.g. Aguirre *et al*., 1998) and still attracts some research interest today.

Adopting a voxel-wise ANCOVA model paved the way for a divergence between the mass-univariate approach used by SPM (i.e. a statistic for each voxel) and multivariate models used previously. A subtle but

Statistical Parametric Mapping, by Karl Friston et al.
ISBN–13: 978-0-12-372560-8 ISBN–10: 0-12-372560-7

important motivation for mass-univariate approaches was the fact that a measured haemodynamic response in one part of the brain may differ from the response in another, *even if the underlying neuronal activation was exactly the same*. This meant that the convention of using region-by-condition interactions as a test for regionally specific effects was not tenable. In other words, even if one showed that two regions activated differently in terms of measured haemodynamics, this did not mean there was a regionally specific difference at the neuronal or computational level. This issue seems to have escaped the electroencephalography (EEG) community, who still use ANOVA with region as a factor, despite the fact that the link between neuronal responses and channel measurements is even more indeterminate than for metabolic imaging. However, the move to voxel-wise, whole-brain analysis entailed two special problems: the problem of registering images from different subjects so that they could be compared on a voxel-by-voxel basis and the multiple-comparisons problem that ensued.

Spatial normalization

The pioneering work of the St Louis group had already established the notion of a common anatomical or stereotactic space (Fox *et al*., 1988) in which to place subtraction or difference maps, using skull X-rays as a reference. The issue was how to get images into that space efficiently. Initially, we tried identifying landmarks in the functional data themselves to drive the registration (Friston *et al*., 1989). This approach was dropped almost immediately because it relied on landmark identification and was not a hundred per cent reproducible. Within a year, a more reliable, if less accurate, solution was devised that matched images to a template without the need for landmarks (Friston *et al*., 1991a). The techniques for spatial normalization using template- or model-based approaches have developed consistently since that time and current treatments regard normalization as the inversion of generative models for anatomical variation that involve warping templates to produce subject-specific images (e.g. Ashburner and Friston, 2005).

Topological inference

Clearly, performing a statistical test at each voxel engendered an enormous false positive rate when using unadjusted thresholds to declare activations significant. The problem was further compounded by the fact that the data were not spatially independent and a simple Bonferroni correction was inappropriate (PET and SPECT (single photon emission computerized tomography) data are inherently very smooth and fMRI had not been invented at this stage). This was the second major theme that occupied people trying to characterize functional neuroimaging data. What was needed was a way of predicting the probabilistic behaviour of SPMs, under the null hypothesis of no activation, which accounted for the smoothness or spatial correlations among voxels. From practical experience, it was obvious that controlling the false positive rate of voxels was not the answer. One could increase the number of positive voxels by simply making the voxels smaller but without changing the topology of the SPM. It became evident that conventional control procedures developed for controlling family-wise error (e.g. the Bonferroni correction) had no role in making inferences on continuous images. What was needed was a new framework in which one could control the false positive rate of the regional effects themselves, noting a regional effect is a topological feature, not a voxel.

The search for a framework for topological inference in neuroimaging started in the theory of stochastic processes and level-crossings (Friston *et al*., 1991b). It quickly transpired that the resulting heuristics were the same as established results from the theory of random fields. Random fields are stochastic processes that conform very nicely to realizations of brain scans under normal situations. Within months, the technology to correct *p*-values was defined within random field theory (Worsley *et al*., 1992). Although the basic principles of topological inference were established at this time, there were further exciting mathematical developments with extensions to different sorts of SPMs and the ability to adjust the *p*-values for small bounded volumes of interest (see Worsley *et al*., 1996). Robert Adler, one of the world's contemporary experts in random field theory, who had abandoned it years before, was understandably very pleased and is currently writing a book with a protégé of Keith Worsley (Adler and Taylor, in preparation).

Statistical parametric mapping

The name 'statistical parametric mapping' was chosen carefully for a number of reasons. First, it acknowledged the TLA of 'significance probability mapping', developed for EEG. Significance probability mapping involved creating interpolated pseudo-maps of *p*-values to disclose the spatiotemporal organization of evoked electrical responses (Duffy *et al*., 1981). The second reason was more colloquial. In PET, many images are derived from the raw data reflecting a number of different physiological parameters (e.g. oxygen metabolism, oxygen extraction fraction, regional cerebral blood flow etc.). These were referred to as parametric maps. All parametric maps are non-linear functions of the original data. The

distinctive thing about *statistical* parametric maps is that they have a known distribution under the null hypothesis. This is because they are predicated on a statistical model of the data (as opposed to a physiological parametric model).

One important controversy, about the statistical models employed, was whether the random fluctuations or error variance was the same from brain region to brain region. We maintained that it was not (on common sense grounds that the frontal operculum and ventricles were not going to show the same fluctuations in blood flow) and adhered to voxel-specific estimates of error. For PET, the Montreal group considered that the differences in variability could be discounted. This allowed them to pool their error variance estimator over voxels to give very sensitive SPMs (under the assumption of stationary error variance). Because the error variance was assumed to be the same everywhere, the resulting *t*-maps were simply scaled subtraction or difference maps (see Fox *et al*., 1988). This issue has not dogged fMRI, where it is generally accepted that error variance is voxel-specific.

The third motivation for the 'statistical parametric mapping' was that it reminded people they were using parametric statistics that assume the errors are additive and Gaussian. This is in contradistinction to non-parametric approaches that are generally less sensitive, more computationally intensive, but do not make any assumptions about the distribution of error terms. Although there are some important applications of non-parametric approaches, they are generally a specialist application in the imaging community. This is largely because brain imaging data conform almost exactly to parametric assumptions by the nature of image reconstruction, post-processing and experimental design.

THE PET YEARS

In the first few years of the nineties, many landmark papers were published using PET and the agenda for a functional neuroimaging programme was established. SPM proved to be the most popular way of characterizing brain activation data. It was encoded in Matlab and used extensively by the MRC Cyclotron Unit at the Hammersmith Hospital in the UK and was then distributed to collaborators and other interested units around the world. The first people outside the Hammersmith group to use SPM were researchers at NIH (National Institutes of Health, UDA) (e.g. Grady *et al*., 1994). Within a couple of years, SPM had become the community standard for analysing PET activation studies and the assumptions behind SPM were largely taken for granted. By this stage, SPM was synonymous with the general linear model and random field theory. Although originally framed in terms of ANCOVA, it was quickly realized that any general linear model could be used to produce an SPM. This spawned a simple taxonomy of experimental designs and their associated statistical models. These were summarized in terms of subtraction or categorical designs, parametric designs and factorial designs (Friston *et al*., 1995a). The adoption of factorial designs was one of the most important advances at this point. The first factorial designs focused on adaptation during motor learning and studies looking at the interaction between a psychological and pharmacological challenge in psychopharmacological studies (e.g. Friston *et al*., 1992). The ability to look at the effect of changes in the level of one factor on activations induced by another led to a rethink of cognitive subtraction and pure insertion and the appreciation of context-sensitive activations in the brain. The latitude afforded by factorial designs is reflected in the fact that most studies are now multifactorial in nature.

THE fMRI YEARS

In 1992, at the annual meeting of the Society of Cerebral Blood Flow and Metabolism in Miami, Florida, Jack Belliveau presented, in the first presentation of the opening session, provisional results using photic stimulation with fMRI. This was quite a shock to the imaging community that was just starting to relax: most of the problems had been resolved, community standards had been established and the way forward seemed clear. It was immediately apparent that this new technology was going to reshape brain mapping radically, the community was going to enlarge and established researchers were going to have to re-skill. The benefits of fMRI were clear, in terms of the ability to take many hundreds of scans within one scanning session and to repeat these sessions indefinitely in the same subject. Some people say that the main advances in a field, following a technological breakthrough, are made within the first few years. Imaging neuroscience must be fairly unique in the biological sciences, in that exactly five years after the inception of PET activation studies, fMRI arrived. The advent of fMRI brought with it a new wave of innovation and enthusiasm.

From the point of view of SPM, there were two problems, one easy and one hard. The first problem was how to model evoked haemodynamic responses in fMRI time-series. This was an easy problem to resolve because SPM could use any general linear model, including convolution models of the way haemodynamic responses were caused (Friston *et al*., 1994). Stimulus functions encoding the occurrence of a particular event or experimental

state (e.g. boxcar-functions) were simply convolved with a haemodynamic response function (HRF) to form regressors in a general linear model (*cf* multiple linear regression).

Serial correlations

The second problem that SPM had to contend with was the fact that successive scans in fMRI time-series were not independent. In PET, each observation was statistically independent of its precedent but, in fMRI coloured time-series, noise rendered this assumption invalid. The existence of temporal correlations originally met with some scepticism, but is now established as an important aspect of fMRI time-series. The SPM community tried a series of heuristic solutions until it arrived at the solution presented in Worsley and Friston (1995). This procedure, also known as 'pre-colouring', replaced the unknown endogenous autocorrelation by imposing a known autocorrelation structure. Inference was based on the Satterthwaite conjecture and is formally identical to the non-specificity correction developed by Geisser and Greenhouse in conventional parametric statistics. An alternative approach was 'pre-whitening' which tried to estimate a filter matrix from the data to de-correlate the errors (Bullmore *et al*., 2001). The issue of serial correlations, and more generally non-sphericity, is still important and attracts much research interest, particularly in the context of maximum likelihood techniques and empirical Bayes (Friston *et al*., 2002).

New problems and old problems

The fMRI community adopted many of the developments from the early days of PET. Among these were the use of the standard anatomical space provided by the atlas of Talairach and Tournoux (1988) and conceptual issues relating to experimental design and interpretation. Many debates that had dogged early PET research were resolved rapidly in fMRI; for example, 'What constitutes a baseline?' This question, which had preoccupied the whole community at the start of PET, appeared to be a non-issue in fMRI with the use of well-controlled experimental paradigms. Other issues, such as global normalization were briefly revisited, given the different nature of global effects in fMRI (multiplicative) relative to PET (additive). However, one issue remained largely ignored by the fMRI community. This was the issue of adjusting *p*-values for the multiplicity of tests performed. While people using SPM quite happily adjusted their *p*-values using random field theory, others seemed unaware of the need to control false positive rates. The literature now entertained reports based on uncorrected *p*-values, an issue which still confounds editorial decisions today. It is interesting to contrast this, historically, with the appearance of the first PET studies.

When people first started reporting PET experiments there was an enormous concern about the rigor and validity of the inferences that were being made. Much of this concern came from outside the imaging community who, understandably, wanted to be convinced that the 'blobs' that they saw in papers (usually *Nature* or *Science*) reflected true activations as opposed to noise. The culture at that time was hostile to capricious reporting and there was a clear message from the broader scientific community that the issue of false positives had to be resolved. This was a primary motivation for developing the machinery to adjust *p*-values to protect against family-wise false positives. In a sense, SPM was a reaction to the clear mandate set by the larger community, to develop a valid and rigorous framework for activation studies. In short, SPM was developed in a culture of scepticism about brain mapping that was most easily articulated by critiquing its validity. This meant that the emphasis was on specificity and reproducibility, as opposed to sensitivity and flexibility. Current standards for reporting brain mapping studies are much more forgiving than they were at its beginning, which may explain why recent developments have focused on sensitivity (e.g. Genovese *et al*., 2002).

The convolution model

In the mid-nineties, there was lots of fMRI research; some of it was novel, some recapitulating earlier findings with PET. From a methodological point of view, notable advances included the development of event-related paradigms that furnished an escape from the constraints imposed by block designs and the use of retinotopic mapping to establish the organization of cortical areas in human visual cortex. This inspired a whole sub-field of cortical surface mapping that is an important endeavour in early sensory neuroimaging. For SPM there were three challenges that needed to be addressed:

Temporal basis functions

The first involved a refinement of the models of evoked responses. The convolution model had become a cornerstone for fMRI with SPM. The only remaining issue was the form of the convolution kernel or haemodynamic response function that should be adopted and whether the form changed from region to region. This was resolved simply by convolving the stimulus function with not one response function but several [basis

functions]. This meant that one could model condition, voxel and subject-specific haemodynamic responses using established approaches. Temporal basis functions (Friston *et al*., 1995b) were important because they allowed one to define a family of HRFs that could change their form from voxel to voxel. Temporal basis functions found an important application in the analysis of event-related fMRI. The general acceptance of the convolution model was consolidated by the influential paper of Boynton a year later (Boynton *et al*., 1996). However, at this time, people were starting to notice some non-linearities in fMRI responses (Vazquez and Noll, 1998) that were formulated, in the context of SPM, as a Volterra series expansion of the stimulus function (Friston *et al*., 1998). This was simple because the Volterra series can be formulated as another linear model (compare with a Taylor expansion). These Volterra characterizations would later be used to link empirical data and balloon models of haemodynamic responses.

Efficiency and experimental design

The second issue that concerned the developers of SPM arose from the growing number and design of event-related fMRI studies. This was the efficiency with which responses could be detected and estimated. Using an analytical formulation, it was simple to show that the boxcar paradigms were much more efficient that event-related paradigms, but event-related paradigms could be made efficient by randomizing the occurrence of particular events such that they 'bunched' together to increase experimental variance. This was an interesting time in the development of data analysis techniques because it enforced a signal processing perspective on the general linear models employed.

Hierarchical models

The third area motivating the development of SPM was especially important in fMRI and reflects the fact that many scans can be obtained in many individuals. Unlike in PET, the within-subject scan-to-scan variability can be very different from the between-subject variability. This difference in variability has meant that inferences about responses in a single subject (using within-subject variability) are distinct from inferences about the population from which that subject was drawn (using between-subject variability). More formally, this distinction is between fixed- and random-effects analyses. This distinction speaks to hierarchical observation models for fMRI data. Because SPM only had the machinery to do single-level (fixed-effects) analyses, a device was required to implement random-effects analyses. This turned out to be relatively easy and intuitive: subject-specific effects were estimated in a first-level analysis and the contrasts of parameter estimates (e.g. activations) were then re-entered into a second-level SPM analysis (Holmes and Friston, 1998). This recursive use of a single-level statistical model is fortuitously equivalent to multilevel hierarchical analyses (compare with the summary statistic approach in conventional statistics).

Bayesian developments

Understanding hierarchical models of fMRI data was important for another reason: these models support empirical Bayesian treatments. Empirical Bayes was one important component of a paradigm shift in SPM from classical inference to a Bayesian perspective. From the late nineties, Bayesian inversion of anatomical models had been a central part of spatial normalization. However, despite early attempts (Holmes and Ford, 1993), the appropriate priors for functional data remained elusive. Hierarchical models provided the answer, in the form of empirical priors that could be evaluated from the data themselves. This evaluation depends on the conditional dependence implicit in hierarchical models and brought previous maximum likelihood schemes into the more general Bayesian framework. In short, the classical schemes SPM had been using were all special cases of hierarchical Bayes (in the same way that the original ANCOVA models for PET were special cases of the general linear models for fMRI). In some instances, this connection was very revealing, for example, the equivalence between classical covariance component estimation using restricted maximum likelihood (i.e. ReML) and the inversion of two-level models with expectation maximization (EM) meant we could use the same techniques used to estimate serial correlations to estimate empirical priors on activations (Friston *et al*., 2002).

The shift to a Bayesian perspective had a number of motivations. The most principled was an appreciation that estimation and inference corresponded to Bayesian inversion of generative models of imaging data. This placed an emphasis on generative or forward models for fMRI that underpinned work on biophysical modelling of haemodynamic responses and, indeed, the framework entailed by dynamic causal modelling (e.g. Friston *et al*., 2003; Penny *et al*., 2004). This reformulation led to more informed spatiotemporal models for fMRI (e.g. Penny *et al*., 2005) that effectively estimate the optimum smoothing by embedding spatial dependencies in a hierarchical model. It is probably no coincidence that these developments coincided with the arrival of the Gatsby Computational Neuroscience Unit next to the Wellcome Department of Imaging Neuroscience. The Gatsby housed several experts in Bayesian inversion and

machine learning and the Wellcome was home to many of the SPM co-authors.

The second motivation for Bayesian treatments of imaging data was to bring the analysis of EEG and fMRI data into the same forum. Source reconstruction in EEG and MEG (magnetoencephalography) is an ill-posed problem that depends explicitly on regularization or priors on the solution. The notion of forward models in EEG-MEG, and their Bayesian inversion had been well established for decades and SPM needed to place fMRI on the same footing.

THE MEG-EEG YEARS

At the turn of the century people had started applying SPM to source reconstructed EEG data (e.g. Bosch-Bayard *et al*., 2001). Although SPM is not used widely for the analysis of EEG-MEG data, over the past five years most of the development work in SPM has focused on this modality. The motivation was to accommodate different modalities (e.g. fMRI-EEG) within the same analytic and anatomical framework. This reflected the growing appreciation that fMRI and EEG could offer complementary constraints on the inversion of generative models. At a deeper level, the focus had shifted from generative models of a particular modality (e.g. convolution models for fMRI) and towards models of neuronal dynamics that could explain any modality. The inversion of these models corresponds to true multimodal fusion and is the aim of recent and current developments within SPM.

In concrete terms, this period saw the application of random field theory to SPMs of evoked and induced responses, highlighting the fact that SPMs can be applied to non-anatomical spaces, such as space-peristimulus-time or time-frequency (e.g. Kilner *et al*., 2005). It has seen the application of hierarchical Bayes to the source reconstruction problem, rendering previous heuristics, like L-curve analysis, redundant (e.g. Phillips *et al*., 2002) and it has seen the extension of dynamic causal modelling to cover evoked responses in EEG-MEG (David *et al*., 2006).

This section is necessarily short because the history of SPM stops here. Despite this, a lot of the material in this book is devoted to biophysical models of neuronal responses that can, in principle, explain any modality. Much of SPM is about the inversion of these models. In what follows, we try to explain the meaning of the more important TLAs entailed by SPM (Table 1-1).

TABLE 1-1 Some common TLAs

TLA Three letter acronym
SPM Statistical parametric map(ping)
GLM General linear model
RFT Random field theory
VBM Voxel-based morphometry
FWE Family-wise error
FDR False discovery rate
IID Independent and identically distributed
MRI Magnetic resonance imaging
PET Positron emission tomography
EEG Electroencephalography
MEG Magnetoencephalography
HRF Haemodynamic response function
IRF Impulse response function
FIR Finite impulse response
ERP Event-related potential
ERF Event-related field
MMN Mis-match negativity
PPI Psychophysiological interaction
DCM Dynamic causal model
SEM Structural equation model
SSM State-space model
MAR Multivariate autoregression
LTI Linear time invariant
PEB Parametric empirical Bayes
DEM Dynamic expectation maximization
GEM Generalized expectation maximization
BEM Boundary-element method
FEM Finite-element method

REFERENCES

Adler RJ, Taylor JE Random fields and geometry. In preparation - To be published by Birkhauser.

Aguirre GK, Zarahn E, D'Esposito M (1998) The inferential impact of global signal covariates in functional neuroimaging analyses. *NeuroImage* **8**: 302–06

Ashburner J, Friston KJ (2005) Unified segmentation. *NeuroImage* **26**: 839–51

Bosch-Bayard J, Valdes-Sosa P, Virues-Alba T *et al*. (2001) 3D statistical parametric mapping of EEG source spectra by means of variable resolution electromagnetic tomography (VARETA). *Clin Electroencephalogr* **32**: 47–61

Boynton GM, Engel SA, Glover GH *et al*. (1996) Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* **16**: 4207–21

Bullmore ET, Long C, Suckling J *et al*. (2001) Colored noise and computational inference in neurophysiological (fMRI) time series analysis: resampling methods in time and wavelet domains. *Hum Brain Mapp* **12**: 61–78

David O, Kiebel SJ, Harrison LM *et al*. (2006). Dynamic causal modeling of evoked responses in EEG and MEG. *NeuroImage*, **30**(4): 1255–7

Duffy FH, Bartels PH, Burchfiel JL (1981) Significance probability mapping: an aid in the topographic analysis of brain electrical activity. *Electroencephalogr Clin Neurophysiol* **51**: 455–62

Fox PT, Mintun MA, Raichle ME *et al*. (1986) Mapping human visual cortex with positron emission tomography. *Nature* **323**: 806–9

Fox PT, Mintun MA, Reiman EM *et al*. (1988) Enhanced detection of focal brain responses using intersubject averaging and change-distribution analysis of subtracted PET images. *J Cereb Blood Flow Metab* **8**: 642–53

Friston KJ, Passingham RE, Nutt JG *et al*. (1989) Localisation in PET images: direct fitting of the intercommissural (AC-PC) line. *J Cereb Blood Flow Metab* **9**: 690–705

Friston KJ, Frith CD, Liddle PF *et al*. (1990) The relationship between global and local changes in PET scans. *J Cereb Blood Flow Metab* **10**: 458–66

Friston KJ, Frith CD, Liddle PF *et al*. (1991a) Plastic transformation of PET images. *J Comput Assist Tomogr* **15**: 634–39

Friston KJ, Frith CD, Liddle PF *et al*. (1991b) Comparing functional (PET) images: the assessment of significant change. *J Cereb Blood Flow Metab* **11**: 690–09

Friston KJ, Frith C, Passingham RE *et al*. (1992) Motor practice and neurophysiological adaptation in the cerebellum: a positron tomography study. *Proc R Soc Lond Series B* **248**: 223–28

Friston KJ, Jezzard PJ, Turner R (1994) Analysis of functional MRI time-series. *Hum Brain Mapp* **1**: 153–71

Friston KJ, Holmes AP, Worsley KJ *et al*. (1995a) Statistical parametric maps in functional imaging: a general linear approach. *Hum Brain Mapp* **2**: 189–210

Friston KJ, Frith CD, Turner R *et al*. (1995b) Characterizing evoked hemodynamics with fMRI. *NeuroImage* **2**: 157–65

Friston KJ, Josephs O, Rees G *et al*. (1998) Nonlinear event-related responses in fMRI. *Magnet Reson Med* **39**: 41–52

Friston KJ, Glaser DE, Henson RN *et al*. (2002) Classical and Bayesian inference in neuroimaging: applications. *NeuroImage* **16**: 484–512

Friston KJ, Harrison L, Penny W (2003) Dynamic causal modelling. *NeuroImage* **19**: 1273–302

Genovese CR, Lazar NA, Nichols T (2002) Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *NeuroImage* **15**: 870–78

Grady CL, Maisog JM, Horwitz B *et al*. (1994) Age-related changes in cortical blood flow activation during visual processing of faces and location. *J Neurosci* **14**:1450–62

Herscovitch P, Markham J, Raichle ME (1983) Brain blood flow measured with intravenous H2(15)O. I. Theory and error analysis. *J Nucl Med* **24**: 782–89

Holmes A, Ford I (1993) A Bayesian approach to significance testing for statistic images from PET. In *Quantification of brain function, tracer kinetics and image analysis in brain PET*, Uemura K, Lassen NA, Jones T *et al*. (eds). *Excerpta Medica, Int. Cong.* Series no. **1030**: 521–34

Holmes AP, Friston KJ (1998) Generalisability, random effects and population inference. *NeuroImage* **7**: S754

Kilner JM, Kiebel SJ, Friston KJ (2005) Applications of random field theory to electrophysiology. *Neurosci Lett* **374**: 174–78

Lauter JL, Herscovitch P, Formby C *et al*. (1985) Tonotopic organization in human auditory cortex revealed by positron emission tomography. *Hear Res* **20**: 199–205

Lueck CJ, Zeki S, Friston KJ *et al*. (1989) The colour centre in the cerebral cortex of man. *Nature* **340**: 386–89

Penny WD, Stephan KE, Mechelli A *et al*. (2004) Comparing dynamic causal models. *NeuroImage* **22**: 1157–72

Penny WD, Trujillo-Barreto NJ, Friston KJ (2005) Bayesian fMRI time series analysis with spatial priors. *NeuroImage* **24**: 350–62

Phillips C, Rugg MD, Friston KJ (2002) Systematic regularization of linear inverse solutions of the EEG source localization problem. *NeuroImage* **17**: 287–301

Talairach P, Tournoux J (1988) *A stereotactic coplanar atlas of the human brain*. Thieme, Stuttgart

Vazquez AL, Noll DC (1998) Nonlinear aspects of the BOLD response in functional MRI. *NeuroImage* **7**: 108–18

Worsley KJ, Evans AC, Marrett S *et al*. (1992) A three-dimensional statistical analysis for CBF activation studies in human brain. *J Cereb Blood Flow Metab* **12**: 900–18

Worsley KJ, Friston KJ (1995) Analysis of fMRI time-series revisited – again. *NeuroImage* **2**: 173–81

Worsley KJ, Marrett S, Neelin P *et al*. (1996) A unified statistical approach of determining significant signals in images of cerebral activation. *Hum Brain Mapp* **4**: 58–73

CHAPTER

2

Statistical parametric mapping

K. Friston

INTRODUCTION

This chapter summarizes the ideas and procedures used in the analysis of brain imaging data. It provides sufficient background to understand the principles of experimental design and data analysis and serves to introduce the main themes covered by subsequent chapters. These chapters have been organized into six parts. The first three parts follow the key stages of analysis: image transformations, modelling, and inference. These parts focus on identifying, and making inferences about, regionally specific effects in the brain. The final three parts address biophysical models of distributed neuronal responses, closing with analyses of functional and effective connectivity.

Characterizing a regionally specific effect rests on estimation and inference. Inferences in neuroimaging may be about differences expressed when comparing one group of subjects to another or, within subjects, changes over a sequence of observations. They may pertain to structural differences (e.g. in voxel-based morphometry) (Ashburner and Friston, 2000) or neurophysiological measures of brain functions (e.g. fMRI or functional magnetic resonance imaging). The principles of data analysis are very similar for all of these applications and constitute the subject of this and subsequent chapters. We will focus on the analysis of fMRI time-series because this covers many of the issues that are encountered in other modalities. Generally, the analyses of structural and PET (positron emission tomography) data are simpler because they do not have to deal with correlated errors from one scan to the next. Conversely, EEG and MEG (electro- and magnetoencephalography) present special problems for model inversion, however, many of the basic principles are shared by fMRI and EEG, because they are both caused by distributed neuronal dynamics. This chapter focuses on the design and analysis of neuroimaging studies. In the next chapter, we will look at conceptual and mathematical models that underpin the operational issues covered here.

Background

Statistical parametric mapping is used to identify regionally specific effects in neuroimaging data and is a prevalent approach to characterizing functional anatomy, specialization and disease-related changes. The complementary perspective, namely functional integration, requires a different set of approaches that examine the relationship among changes in one brain region relative to changes in others. Statistical parametric mapping is a voxel-based approach, employing topological inference, to make some comment about regionally specific responses to experimental factors. In order to assign an observed response to a particular brain structure, or cortical area, the data are usually mapped into an anatomical space. Before considering statistical modelling, we deal briefly with how images are realigned and normalized into some standard anatomical space. The general ideas behind statistical parametric mapping are then described and illustrated with attention to the different sorts of inferences that can be made with different experimental designs.

EEG, MEG and fMRI data lend themselves to a signal processing perspective. This can be exploited to ensure that both the design and analysis are as efficient as possible. Linear time invariant models provide the bridge between inferential models employed by statistical mapping and conventional signal processing approaches. We will touch on these and develop them further in the next chapter. Temporal autocorrelations in noise processes represent another important issue, especially in fMRI, and approaches to maximizing efficiency in the context of serially correlated errors will be discussed. We will also consider event and epoch-related designs in terms of efficiency. The chapter closes by looking at the distinction

Statistical Parametric Mapping, by Karl Friston et al.
ISBN–13: 978-0-12-372560-8 ISBN–10: 0-12-372560-7

between fixed and random-effect analyses and how this relates to inferences about the subjects studied or the population from which these subjects came.

In summary, this chapter reviews the three main stages of data analysis: spatial or image transforms, modelling and inference; these are the areas covered in the first three parts of this book and are summarized schematically in Plate 1 (see colour plate section). We then look at experimental design in light of the models covered in earlier parts. The next chapter deals with different models of distributed responses and previews the material covered in the final three parts of this book.

SPATIAL TRANSFORMS AND COMPUTATIONAL ANATOMY

A central theme in this book is the inversion of forward or generative models of how data are caused. We will see this in many different contexts, from the inversion of linear models of fMRI time-series to the inversion of dynamic causal models of distributed EEG responses. Image reconstruction, in imaging modalities like PET and fMRI, can be regarded as inverting a forward model of how signals, deployed in anatomical space, conspire to produce measured signals. In other modalities, like EEG and MEG, this inversion, or source reconstruction, can be a substantial problem in its own right. In most instances, it is expedient to decompose the inversion of forward spatiotemporal models into spatial and temporal parts. Operationally, this corresponds to reconstructing the spatial signal at each time point and then inverting a temporal model of the time-series at each spatial source (although we will consider full spatiotemporal models in Chapters 25 and 26). This view of source or image reconstruction as model inversion can be extended to cover the inversion of anatomical models describing anatomical variation within and between subjects. The inversion of these models corresponds to registration and normalization respectively. The aim of these anatomical inversions or transformations is to remove or characterize anatomical differences. Chapters 4 to 6 deal with the inversion of anatomical models for imaging modalities. Figure 2.1 shows an example of a generative model for structural images that is presented in Chapter 6. Chapters 28 and 29 deal with the corresponding inversion for EEG and MEG data.

This inversion corresponds to a series of spatial transformations that try to reduce unwanted variance components in the voxel time-series. These components are induced by movement or shape differences among a series of scans. Voxel-based analyses assume that data from a particular voxel derive from the same part of the brain. Violations of this assumption will introduce artefactual changes in the time-series that may obscure changes, or differences, of interest. Even single-subject analyses usually proceed in a standard anatomical space, simply to enable reporting of regionally-specific effects in a frame of reference that can be related to other studies. The first step is to realign the data to undo the effects of subject movement during the scanning session (see Chapter 4). After realignment, the data are then transformed using linear or non-linear warps into a standard anatomical space (see Chapters 5 and 6). Finally, the data are usually spatially smoothed before inverting the temporal part of the model.

Realignment

Changes in signal intensity over time, from any one voxel, can arise from head motion and this represents a serious confound, particularly in fMRI studies. Despite restraints on head movement, cooperative subjects still show displacements of up several millimetres. Realignment involves estimating the six parameters of an affine 'rigid-body' transformation that minimizes the differences between each successive scan and a reference scan (usually the first or the average of all scans in the time series). The transformation is then applied by resampling the data using an interpolation scheme. Estimation of the affine transformation is usually effected with a first-order approximation of the Taylor expansion of the effect of movement on signal intensity using the spatial derivatives of the images (see below). This allows for a simple iterative least square solution that corresponds to a Gauss-Newton search (Friston *et al*., 1995a). For most imaging modalities this procedure is sufficient to realign scans to, in some instances, a hundred microns or so (Friston *et al*., 1996a). However, in fMRI, even after perfect realignment, movement-related signals can still persist. This calls for a further step in which the data are *adjusted* for residual movement-related effects.

Adjusting for movement-related effects

In extreme cases, as much as 90 per cent of the variance in fMRI time-series can be accounted for by the effects of movement *after* realignment (Friston *et al*., 1996a). Causes of these movement-related components are due to movement effects that cannot be modelled using a linear model. These non-linear effects include: subject movement between slice acquisition, interpolation artefacts (Grootoonk *et al*., 2000), non-linear distortion due to magnetic field inhomogeneities (Andersson *et al*., 2001) and spin-excitation history effects (Friston *et al*., 1996a). The

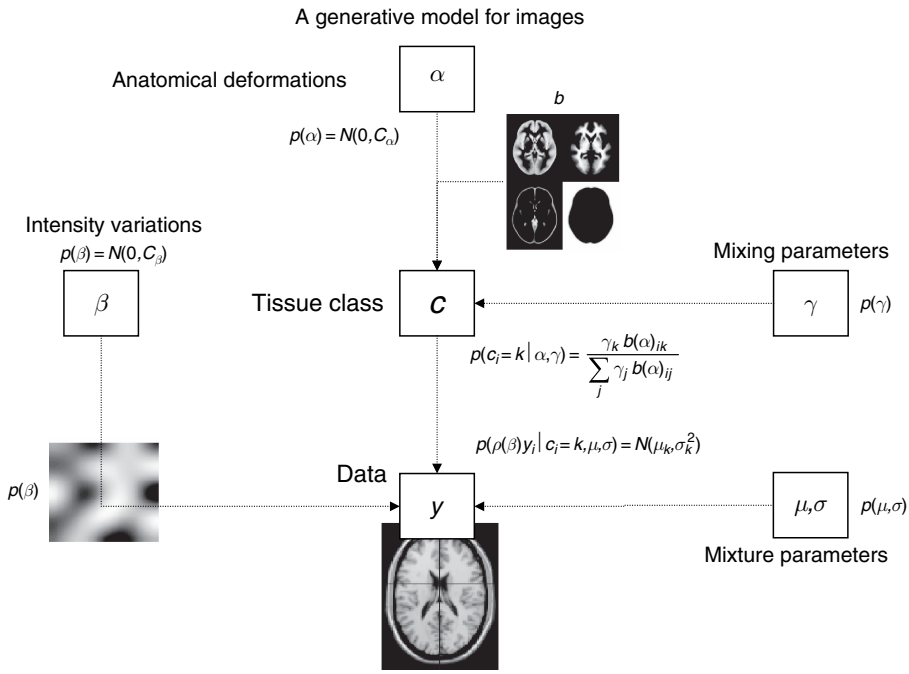

FIGURE 2.1 A graphical model describing the generation of an image. The boxes or 'nodes' represent quantities required to generate an image and the lines or 'edges' encode conditional dependencies among these quantities. This graphical description is a useful way to describe a generative model and makes all the conditional dependencies explicit. In this example, one starts by sampling some warping parameters α from their prior density $p(\alpha)$. These are used to resample (i.e. warp) a series of tissue-class maps to give $b(\alpha)_{ik}$ for each voxel and tissue class. The warping parameters model subject-specific anatomical deformations. Mixing parameters γ are then selected from their prior density $p(\gamma)$; these control the relative proportions of different tissue-classes over the brain. The mixing parameters scale the tissue-class maps to provide a density from which a voxel-specific tissue-class c_i is sampled. This specifies a mixture of Gaussians from which the voxel intensity is sampled. This mixture is specified in terms of the expectations μ and variances σ of their constituent Gaussians that are sampled from the prior density $p(\mu, \sigma)$. The final stage of image construction is to scale the voxel values with some slowly varying intensity field whose parameters β are sampled from their prior $p(\beta)$. The resulting image embodies random effects expressed at the level of anatomical deformation, amount of different tissue types, the expression of those tissues in the measurement, and image-specific inhomogeneities. Inversion of this generative model implicitly corrects for intensity variations, classifies each voxel probabilistically (i.e. segments) and spatially normalizes the image. Critically, this inversion accounts properly for all the conditional dependencies among the model's parameters and provides the most likely estimates given the data (see Chapter 6 for details of this model and its inversion).

latter can be pronounced if the repetition time approaches T1 making the current signal a function of movement history. These effects can render the movement-related signal a non-linear function of displacement in the n-th and previous scans:

$$y_n = f(x_n, x_{n-1}, \ldots)$$

By assuming a sensible form for this function, one can include these effects in the temporal model, so that they are explained away when making inferences about activations. This relies on accurate displacement estimates from the realignment and assumes activations are not correlated with the movements (any component that is correlated will be explained away).

The form for $f(x_n, x_{n-1}, \ldots)$, proposed in Friston *et al.* (1996a), was a non-linear autoregression model that used polynomial expansions to second order. This model was motivated by spin-excitation history effects and allowed displacement in previous scans to explain movement-related signal in the current scan. However, it is also a reasonable model for other sources of movement-related confounds. Generally, for repetition times (TR) of several seconds, interpolation artefacts supersede (Grootoonk *et al.*, 2000) and first-order terms, comprising an expansion of the current displacement in terms of periodic basis functions, are sufficient.

This section has considered *spatial* realignment. In multislice acquisition, different slices are acquired at different times. This raises the possibility of *temporal* realignment to ensure that the data from any given volume were sampled at the same time. This is usually performed using interpolation over time and only when the TR is sufficiently small to permit interpolation. Generally, timing effects of this sort are not considered problematic because they manifest as artefactual latency differences in evoked responses from region to region. Given that biological latency differences are in the order of a few seconds, inferences about