Neuropsychologia xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

Neuropsychologia

journal homepage: www.elsevier.com/locate/neuropsychologia

The dimensionalities of lesion-deficit mapping

Tianbo Xu[a], Ashwani Jha[a,b], Parashkev Nachev[a,b,*]

[a] Institute of Neurology, UCL, UK
[b] National Hospital for Neurology and Neurosurgery, Queen Square, UK

ARTICLE INFO

Keywords:
Lesion-deficit brain mapping
Structural neuroimaging
Multivariate analysis
Machine learning
Stroke

ABSTRACT

Lesion-deficit mapping remains the most powerful method for localising function in the human brain. As the highest court of appeal where competing theories of cerebral function conflict, it ought to be held to the most stringent inferential standards. Though at first sight elegantly transferable, the mass-univariate statistical framework popularized by functional imaging is demonstrably ill-suited to the task, both theoretically and empirically. The critical difficulty lies with the handling of the data's intrinsically high dimensionality. Conceptual opacity and computational complexity lead lesion-deficit mappers to neglect two distinct sets of anatomical interactions: those between areas unified by function, and those between areas unified by the natural pattern of pathological damage. Though both are soluble through high-dimensional multivariate analysis, the consequences of ignoring them are radically different. The former will bleach and coarsen a picture of the functional anatomy that is nonetheless broadly faithful to reality; the latter may alter it beyond all recognition. That the field continues to cling to mass-univariate methods suggests the latter problem is misidentified with the former, and that their distinction is in need of elaboration. We further argue that the vicious effects of lesion-driven interactions are not limited to anatomical localisation but will inevitably degrade purely predictive models of function such as those conceived for clinical prognostic use. Finally, we suggest there is a great deal to be learnt about lesion-mapping by simulation-based modelling of lesion data, for the fundamental problems lie upstream of the experimental data themselves.

1. Introduction

In common with all scientific inference, the fidelity of lesion-deficit mapping depends on the quality of the source data and the validity of the models applied to it. Though equally important, the two aspects are sharply distinct: a deficit in neither is remediable by an excess of the other. Whereas a good model may be improved by better data, a defective model is often irredeemably so. The validity of a model is judged by hard, logico-mathematical criteria, the quality of data by softer, empirical opinion. Inferential failure resulting from poor data tends to be graceful, proportionate with the degree of data corruption; by contrast, model errors may have catastrophic consequences even when seemingly minor. Worse, failure from a defective model is often *silent*, cloaked in superficially attractive significance values that conceal fatal *biases* in the inference repetition can only entrench. Where no other inferential technique is stronger, such systematic errors may easily persist indefinitely.

Why do we need reminding of these statistical platitudes? The hazards of modelling are greatest where the complexity of the system under study is highest, as is archetypally true of the brain. For our purposes it suffices to define complexity as the minimum number of dimensions required to predict one state of a system from another: its *intrinsic dimensionality*. If our models cannot be commensurately complex—for reasons of intellectual opacity or computational tractability—it is tempting aggressively to simplify them, for then the unmodelled signal superficially resembles noise. But if the residual variance retains appreciable structure, the inference will be distorted in ways the simplicity of the model merely conceals from view. The inevitable inferential distortion aside, the more non-stochastic variability the model does not explain, the weaker its explanatory power, and—of course—its practical, clinical utility.

So how do we determine the correct dimensionality? A perfect answer is impossible, for it assumes precisely the knowledge our models are deployed to acquire. But we can examine the grounds for an informed supposition, and we can also explicitly *test* the consequences of adopting one solution over another. Here we give the empirical and conceptual grounds for our view on the necessary dimensionality, and go on to outline the explicit tests one ought to conduct to confirm or infirm it. Although this is certainly not the only important methodological concern in lesion-deficit mapping, we dwell on it at length here

* Corresponding author at: Institute of Neurology, UCL, UK.
E-mail address: p.nachev@ucl.ac.uk (P. Nachev).

http://dx.doi.org/10.1016/j.neuropsychologia.2017.09.007
Received 31 May 2017; Received in revised form 9 August 2017; Accepted 7 September 2017

because it has received so little of the attention it requires.

2. The dimensionality of anatomical inference in the brain

Let us first clarify how the dimensionality of an inferential model is determined. Formal lesion-deficit mapping began with taking the overlap of a set of lesions and contrasting the peak with that derived from another, control set of patients (e.g. (Robertson et al., 1988)). Such a comparison is produced by a simple voxel-wise operation that ignores any anatomical relationship but that between homologous voxels across the two groups. This is mass-univariate inference, even if it was not called so at the time, for the contribution of each voxel is independently quantified by its own, univariate test, whether implicit or explicit. By replacing simple subtraction with a formal statistical test, voxel-wise lesion-symptom mapping (VLSM) and kindred techniques add a measure of confidence to the inference at each voxel, leaving the independence assumption untouched, and the inference univariate (Bates et al., 2003; Chen and Herskovits, 2010; Damasio et al., 1996; Karnath et al., 2004; Rorden et al., 2007). Additional variables may be added to the voxel-wise statistical test—various behavioural covariates, for example—making it multivariate, but *not* from the critical perspective of the anatomy, for that is still modelled as a set of independent locations, evaluated over multiple statistical tests run at each voxel in isolation from every other. So this is still mass-univariate *anatomical* inference, even if its *behavioural* dimensionality may be expanded.

Now two manoeuvres here commonly escalate the anatomical dimensionality. The most common is the addition of lesion volume as a covariate, a crude index of damage at other voxels (Karnath et al., 2004). This attempts to capture the effect on behaviour of the global change in available brain substrate, independently of anatomical location, reasoning that parcelling out such anatomically non-specific effects will increase sensitivity for the anatomically specific effects of interest. Less common is the use of Gaussian smoothing, which changes the value of a voxel on the assumption is relation to its neighbours is adequately described by a random Gaussian field (Kimberg et al., 2007). Since neither is capable of conveying any substantial anatomical detail, most would still regard such models as mass-univariate. Moreover, we still have one model per voxel, and therefore as many models as there are estimated voxels in the brain.

An analysis becomes anatomically multivariate where the statistical model incorporates many anatomical variables, indexing the presence or absence of damage to different parts of the brain *together* (Chen et al., 2008; Chen and Herskovits, 2015; Keinan et al., 2004; Mah et al., 2014; Rondina et al., 2016; Smith et al., 2013; Toba et al., 2017; Yourganov et al., 2016; Zhang et al., 2014). The dimensionality of such models depends on the number of such variables and their properties. Where the variables are correlated, the intrinsic dimensionality will be less than their number, but this is usually something to be established by the analysis itself, implicitly in the inferential model, or explicitly in a preceding dimensionality reduction step. Either way, each inferential model now covers all or a substantial part of the brain, leaving us with one or few models per brain where a multiplicity of voxels describe a large number of dimensions per model.

Naturally, the dimensions of behaviour and anatomy are bound to interact, and a model may be critically deficient in either or both. Our focus here is on the anatomical not because the others should be neglected but because the anatomical near-universally have.

3. Two determinants of dimensionality: brain and lesions

It is natural to think of anatomical factors as pertaining only to the functional architecture of the brain. But in lesion-deficit mapping this is only one side of the coin: there is a second anatomical dimensionality to consider, that arising from the lesion architecture. We need to examine each in turn.

3.1. Brain dimensionality

That Lego® is not helpfully metaphorical of the brain's functional architecture is increasingly recognized in the emphasis on highly distributed operations subserved by complex, dynamic functional networks (Sporns et al., 2005). Both disruptive and correlative data unequivocally point to an underlying neural organisation in which complex *interactions* between areas determine the observed behaviour (Young et al., 2000). Such interactions may be non-monotonic, reflective of neural relations that could just as easily be competitive as collaborative. They are—moreover—bound to be *adaptive*, varying across both time and individuals. An entire field of clinical neuroscience—functional neurosurgery—richly illustrates these truths in each and every patient, where *disruption* of one area of the brain—optimised both within and across patients—is used to improve the function of the brain as a whole (Jha and Brown, 2011; Johnson et al., 2008).

A satisfactory model of a lesioned brain must therefore not only model the individual functions of the affected areas but their—potentially highly complex—interactions. The syndrome of visuospatial neglect offers a striking example of this: neglect caused by damage to inferior parietal areas may not only *not* be exacerbated by damage to the contralateral frontal eye field but wholly *reversed* by it (Vuilleumier et al., 1996). It is obvious that in evaluating the lesion-deficit relationship in a patient we must here model the presence and absence of damage at *both* loci, *together*, and if this is true of this particular pair it may be true of any combination of areas, across the entire brain (Price and Friston, 2002; Zavaglia and Hilgetag, 2016).

The optimal lesion-deficit model, then, is one in which the integrity of each functionally homogeneous location in the brain is a separate variable. Since no wholly convincing definition of functional homogeneity is currently available (pace (Glasser et al., 2016)), our limit becomes practical: such anatomical parcellation of the brain as our tools can provide, minimally the voxel size of the imaging acquisition. Anything short of this will miss interactions at a finer level of anatomical organisation. Even with voxel sizes of remarkable coarseness—8 mm isotropic—this leaves us with several thousand variables per brain: a high-dimensional model, certainly in proportion to the number of patients included in the typical lesion-deficit study.

3.2. Lesion dimensionality

The variable expansion we are discussing here is driven by the dimensionality of the functional architecture. But in lesion-deficit mapping there is a second, independent dimensionality to consider: that of the lesion architecture (Mah et al., 2014, 2015). Where lesions overlap—and are generally larger than the minimal size of functionally homogeneous areas—the lesion-deficit relation will be influenced by both the functional and the lesion architecture. This is overwhelmingly true of the lesions described in the current literature, a reflection of the natural characteristics of the underlying pathology, especially the commonest: vascular injury.

Let us consider carefully why the lesion architecture matters here. In functional imaging, the physiological cause of the change in the BOLD signal operates at sub-voxel granularity, for it is driven by the microvasculature (Logothetis et al., 2001). Such anatomical structure as emerges at the voxel level is then plausibly related to the underlying neural anatomy, even if there may well be non-linearities in the relation between BOLD and neural activity across the brain (Birn et al., 2001; Heeger and Ress, 2002). If two voxels are co-activated it will *not* be because an idiosyncrasy of the microvasculature makes it so, for the vascular causal mechanisms do not operate at that anatomical scale. Consider, by contrast, lesion-deficit mapping, where the effective equivalent of BOLD activation is a lesion, almost invariably extending across multiple voxels as an outcome *not* of the underlying functional anatomy but of the causal pathological process. The anatomy of the