Decoding the anatomical network of spatial attention

David V. Smith[a,b,1], John A. Clithero[b,2], Christopher Rorden[c], and Hans-Otto Karnath[c,d,1]

[a]Department of Psychology and Neuroscience and [b]Center for Cognitive Neuroscience, Duke University, Durham, NC 27708; [c]Department of Psychology, University of South Carolina, Columbia, SC, 29208; and [d]Center of Neurology, Division of Neuropsychology, Hertie-Institute for Clinical Brain Research, University of Tuebingen, 72076 Tuebingen, Germany

The study of stroke patients with modern lesion-symptom analysis techniques has yielded valuable insights into the representation of spatial attention in the human brain. Here we introduce an approach—multivariate pattern analysis—that no longer assumes independent contributions of brain regions but rather quantifies the joint contribution of multiple brain regions in determining behavior. In a large sample of stroke patients, we found patterns of damage more predictive of spatial neglect than the best-performing single voxel. In addition, modeling multiple brain regions—those that are frequently damaged and, importantly, spared—provided more predictive information than modeling single regions. Interestingly, we also found that the superior temporal gyrus demonstrated a consistent ability to improve classifier performance when added to other regions, implying uniquely predictive information. In sharp contrast, classifier performance for both the angular gyrus and insular cortex was reliably enhanced by the addition of other brain regions, suggesting these regions lack independent predictive information for spatial neglect. Our findings highlight the utility of multivariate pattern analysis in lesion mapping, furnishing neuroscience with a modern approach for using lesion data to study human brain function.

brain injury | superior temporal cortex | voxelwise lesion symptom mapping | distributed network

Observing the behavioral consequences of brain injury has driven our understanding of brain function. Although recent brain activation measures have revealed how behavior engages spatially distributed networks, lesion methods remain important because of their clinical significance and level of inference, as a result of their ability to detect if a region is critical for a task rather than merely involved with a task (1, 2). Classically, lesion–behavior relationships were inferred by looking for associations: identification of regions consistently damaged in patients with a given symptom. However, this approach suffers from a major confound; some regions of the brain are more likely to be injured than others, and therefore, these studies identify both regions that are critical to a function as well as regions that are frequently injured (for review, see ref. 2). Voxelwise lesion symptom mapping (VLSM) revolutionized this method by looking for statistical dissociations: identifying regions that are consistently damaged in individuals with a deficit but spared in those without a deficit (3, 4).

Conventional VLSM methods compute independent analyses for each and every voxel of the brain. Unfortunately, this mass univariate method necessarily limits the statistical power for many common neurological syndromes. For example, if damage to either of two distant brain regions can lead to the same symptoms (e.g., a distributed network), a patient with damage to one location will effectively appear as a counter example for detecting the other location. Indeed, consider the situation where a symptom is observed whenever only a portion of a large functional module is injured—in this case, two patients who have the same symptoms and damage to the same functional module may have mutually exclusive regions of injury and thus appear to provide statistical counter examples for each other (the partial injury problem, see ref. 4).

Multivariate pattern analysis (MVPA) can address these shortcomings. Although MVPA has yielded remarkable insights into brain–behavior relationships in cognitive neuroscience (5, 6), its techniques have not been applied to traditional lesion mapping approaches. Analyzing brain injury data using MVPA—a method for simultaneously considering the influence of multiple regions or voxels—could extend our understanding of both brain function and stark behavioral deficits incurred by brain damage. Beyond leveraging spatially distributed regions critical for a task, multivariate methods are also sensitive to regions where injury predicts normal performance on a specific task. By combining (noisy) information of regions that predict both the presence of impaired performance and the retention of performance after injury, a classifier should become more accurate (7). These strengths suggest that multivariate methods would provide improved sensitivity over univariate methods for predicting behavioral performance based on brain imaging acquired following brain injury. In particular, MVPA is well suited to symptoms that are observed following damage to portions of a large module or spatially distributed networks, as in spatial neglect (8, 9).

The present study uses MVPA to characterize how multiple right hemisphere brain regions contribute to predicting spatial neglect in a large sample of right-sided stroke patients. Individuals with neglect show profound impairments, including spontaneous deviation of eyes and head toward the ipsilesional side (10, 11) and ignoring objects located on the contralesional side (12–15). A network of brain areas—many of which are commonly implicated in the deployment and allocation of spatial attention (16–18)—appear to be involved with this disorder (for review, see refs. 8 and 9), including the superior and middle temporal cortex and underlying insula (19–23), the inferior parietal cortex and temporo-parietal junction area (22–25), as well as the inferior frontal cortex (23, 26). Further, spatial neglect is observed following injury to the fiber tracts interconnecting these cortical sites (27–30).

Given the diverse brain regions implicated in spatial neglect, there are clear and distinct hypotheses regarding the contribution of different anatomical nodes in predicting this disorder. In this situation, it is obvious that a multivariate approach—one that leverages information from multiple brain regions, including those that are commonly spared and damaged—could help to better understand the neural bases of this disorder. We therefore used two machine-learning techniques, linear and nonlinear support vector machines (SVMs; a toy linear example to illustrate our method is displayed in Fig. 1A), to classify individuals based on structural brain scans with right hemisphere lesions demarcated in standard space (Fig. 1B). Predictive performance for each
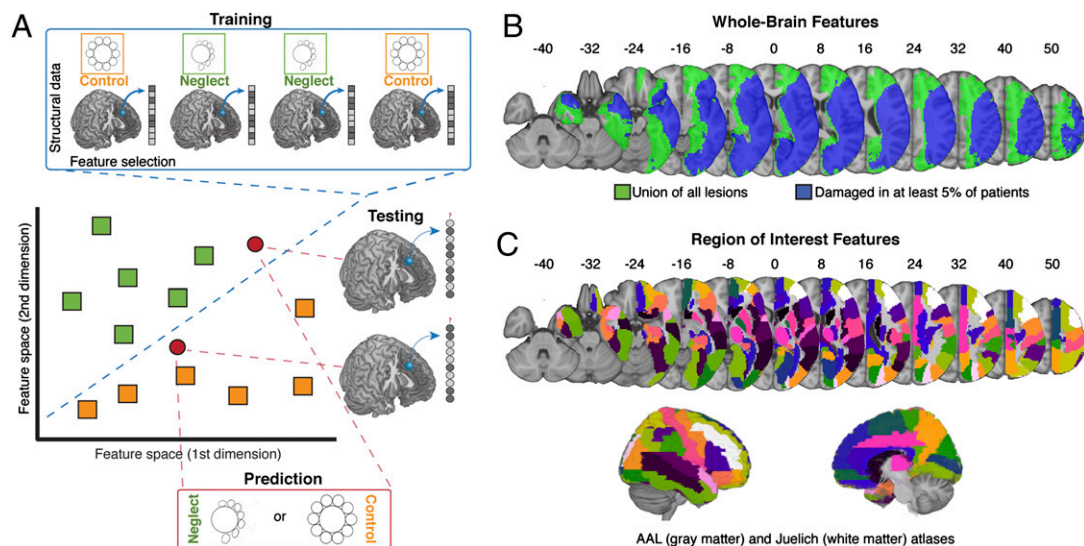
Fig. 1. Multivariate analytic approach. (*A*) We used a MVPA procedure for using brain injury maps to predict the presence or absence of spatial neglect. MVPA involves both training and testing a predictive model. The training procedure used machine-learning algorithms [support vector machines (SVMs)] to construct a model of how distributed patterns of lesion data indicate the presence or absence of spatial neglect. The constructed model is then tested on new data (i.e., not used to train the model). This procedure was repeated for each individual in our dataset ($n = 140$), meaning each individual was tested on a model that was built independently of that individual's data. The average of those predictions is the predictive power of the model. The algorithm can be trained on various features—in this case voxels from different regions of the brain—to predict the presence or absence of neglect. The feature space is typically very high dimensional, as many voxels are available, but the figure displays a simple 2D case for illustrative purposes. (*B*) Feature spaces used for our whole-brain analyses included both the set of voxels equal to the union of all individual right-hemisphere lesions (green) and voxels containing lesions found in at least 5% of the subject population (blue). (*C*) We also generated more selective feature spaces ($n = 56$) using regions of interest from the AAL (gray matter) and Juelich (white matter) atlases. Axial slice numbers are provided in terms of MNI space.

classifier was based on a leave-one-out cross-validation (CV) testing procedure (*SI Materials and Methods*). The cross-validation testing procedure provides a robust measure for how well available data can be used to make individual out-of-sample predictions for the presence or absence of spatial neglect. We also used a combinatoric analytic approach (31–33), in which multiple brain regions of interest (ROIs; Fig. 1*C*; *SI Materials and Methods*) were iteratively combined, affording us an opportunity to quantify the ability of each ROI to uniquely convey information about the presence or absence of spatial neglect (*SI Materials and Methods*). Our analyses focused on three questions. First, do multivariate techniques outperform traditional univariate techniques? Second, which set of right hemisphere brain regions is most predictive of spatial neglect? Finally, do any regions carry significantly more unique information than other brain regions in the right hemisphere?

Results

Because one of our principal goals was to demonstrate that our multivariate approach offers an advantage over traditional univariate techniques, we initially sought to establish predictive baselines that would be found by conventional analyses using single voxels in isolation. Next, we evaluated whether conventional analyses are outperformed by multivariate classifications that use multiple ROIs. We emphasize that, when using multiple ROIs, counterintuitive results are possible, as multivariate analyses can leverage information from both positive (i.e., regions generally associated with neglect) and negative predictors (i.e., regions not generally associated with neglect). Finally, we focused on 12 right hemisphere regions classically associated with spatial neglect in the literature and used a combinatoric analytic approach (31–33) to examine which of these 12 canonical regions contain unique information regarding the prediction of spatial neglect.

Establishing Predictive Baselines with Univariate Classifications. We first sought to identify single voxels that carried significant

information regarding neglect. To ensure a fair comparison with our core analyses, we used identical machine-learning algorithms, including linear and nonlinear SVMs, and out-of-sample prediction as an index of performance. We found a focal pattern of voxels within the right superior temporal gyrus (STG) that were above chance performance at an uncorrected threshold (Fig. S1*A*). The voxel [Montreal Neurological Institute (MNI)$_{(x,y,z)}$ = 70,−16,6] was significantly greater than chance with both classifiers ($CV_{linear} = 74.57\%$, $P_{perm} < 0.0001$; $CV_{nonlinear} = 73.57\%$, $P_{perm} < 0.0001$). Using this STG voxel, we found that the linear SVM significantly outperformed the nonlinear SVM ($CV_{increase} = 1$, $P_{perm} < 0.002$), presumably reflecting modest overfitting when a nonlinear solution is applied to a linear space (34).

As a control analysis, we also examined how well lesion size could classify spatial neglect and control patients. Lesion size is a confounding factor in all lesion-mapping studies: larger lesions are more likely to compromise critical modules. Consistent with this logic, our results indicated that lesion size was able to significantly classify neglect and control patients above chance for both SVM classifiers ($CV_{linear} = 59.71\%$, $P_{perm} < 0.004$; $CV_{nonlinear} = 78.14\%$, $P_{perm} < 0.0001$). In this case, the nonlinear SVM classifier significantly outperformed the linear SVM classifier ($CV_{increase} = 18.43$, $P_{perm} < 0.0001$), likely reflecting the nonlinear nature of the 1D feature space characterizing lesion size. Owing to the enhanced predictive power of the nonlinear classifier in this baseline test, all subsequent analyses use the nonlinear classifier.

Multivariate Classifications Outperform Univariate Classifications. We predicted that multivariate analyses would outperform traditional univariate analyses when classifying spatial neglect. To examine this prediction, we constructed several distinct sets of feature spaces that included multiple pieces of information for each classification: single voxels with lesion size included as its own feature (*SI Results*; Fig. S1*B*); whole-brain (i.e., all usable features across brain regions from our subject population) feature spaces with and without lesion size included as its own

Smith et al.

feature (*SI Results*); and ROIs with and without lesion size included as its own feature.

Here, we focus on the ROI-based analyses, as those feature spaces are less likely to contain gross imbalances between informative and uninformative features. We examined 45 gray and white matter regions taken from the Automated Anatomical Labeling (AAL) and Juelich white matter atlases (Fig. 1*C*; Table S1) and performed independent classifications on each. We found several regions, including the STG, whose performance was significantly above chance, even after correcting for multiple comparisons (Table S2). We also assessed whether including lesion size as its own feature influenced the performance of any of our 45 regions. Across these regions, we identified several regions in visual cortex and prefrontal cortex whose performance increased significantly as a result of modeling lesion size as a separate feature (Table S3 and Fig. S2). Importantly, none of these regions were identified as having above-chance predictive power without including lesion size as its own feature, indicating a statistical dependency on lesion size within these regions.

After identifying single brain regions whose pattern of damage could distinguish neglect and control patients, we next evaluated a key prediction: that additional regions could significantly improve predictive power above and beyond what could be expected by examining single regions in isolation. To test this prediction, we created two additional sets of feature spaces that were formed using all combinations of two and three ROIs (990 combinations of two ROIs; 14,190 combinations of three ROIs). We first compared the average CV of all single ROIs, combinations of two ROIs, and combinations of three ROIs (Fig. 2*A*). As the underlying distribution corresponding to changes in average CV percentage is unknown, we constructed null distributions using Monte Carlo permutation-based testing (*SI Materials and Methods*). This method entailed running each combinatoric analysis ($n = 15,225$) 1,000 times with permuted neglect and control labels. For each permutation, we computed the change in average CV percentage for both comparisons (i.e., 2 ROIs – 1 ROI and 3 ROIs – 2 ROIs). Our results showed that the average CV percentage for combinations of two ROIs was significantly higher compared with single ROIs ($CV_{increase} = 6.3$, $P < 0.001$; Fig. 2*B*). Additionally, we found that the average CV percentage was significantly higher for combinations of three ROIs compared with combinations of two ROIs ($CV_{increase} = 3.1$, $P < 0.001$; Fig. 2*C*). Our results also survived an important robustness check: a similar pattern of results were identified after regressing out the confounding influence of ROI size and the disparity in ROI damage between neglect and control patients (Fig. S3; Table S4).

As a conservative and principled set of tests, we next used the best performing single ROI, corticospinal tract (CV = 75.85%), as a baseline for our comparisons against combinations of two and three ROIs. We found no combination of two ROIs that outperformed the best single ROI. Strikingly, we found that one combination of three ROIs—middle occipital gyrus, Heschl's gyrus, and corticospinal tract—significantly outperformed the best single ROI ($CV_{increase} = 10.42$, $P_{perm} < 0.03$) for an overall CV of 86.27%. These three regions, which are not generally associated with spatial neglect (8, 9), highlight one of the significant advances of multivariate classification techniques: the ability to capitalize on both positive and negative predictors. Nevertheless, interpreting such results presents a challenge, because our metrics for classifier performance do not provide insight into predictive directionality of the individual features that comprise each analysis.

We therefore focused our next analyses on only those 12 critical perisylvian regions in the right hemisphere that have been previously associated with spatial neglect in the literature (Table S1); these analyses evaluate whether a given region significantly improves the performance when added to another region (or set of regions). We found one combination of two ROIs whose
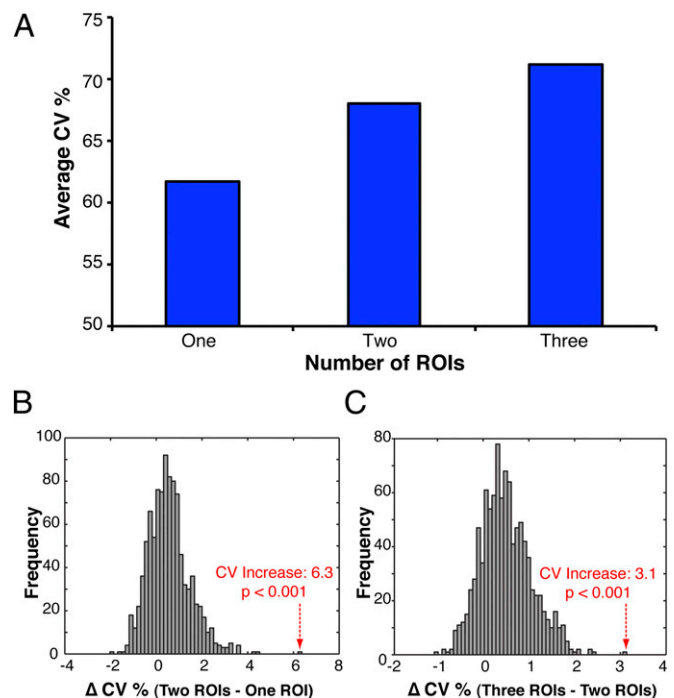

Fig. 2. Modeling multiple brain regions improves predictive power. (*A*) To evaluate whether including more regions provides additional predictive power, we show the average CV percentages across all single ROIs ($n = 45$), all combinations of two ROIs ($n = 990$), and all combinations of three ROIs ($n = 14,190$). Because the underlying distribution corresponding to changes in average CV percentage is unknown, we constructed null distributions running all analyses 1,000 times with permuted neglect and control labels. For each permutation, we computed the change in average CV percentage for both comparisons (i.e., 2 ROIs – 1 ROI; and 3 ROIs – 2 ROIs). (*B*) In comparing the true CV to the null distribution, we found that the average CV significantly increased when adding one ROI to another ROI ($CV_{increase} = 6.3$, $P < 0.001$; red arrow). (*C*) Additionally, in a similar statistical test, we found that the average CV significantly increased when adding one ROI to two ROIs ($CV_{increase} = 3.1$, $P < 0.001$; red arrow). Data in histograms are partitioned into 50 equally spaced bins on the *x*-axis.

performance was significantly better than either region in isolation. Specifically, adding the right STGs to the inferior frontal gyrus pars orbitalis (IF Orb; $CV_{original} = 52.14\%$) significantly improved performance ($CV_{combo2} = 73.57\%$; $CV_{increase} = 21.42$; $P < 0.05$, Bonferroni corrected). We extended this approach and identified two combinations of three ROIs where the addition of the third ROI was significantly better than the two ROIs together. First, adding the right middle temporal pole to the combination of inferior frontal gyrus pars triangularis (IF Tri) and supramarginal gyrus ($CV_{combo2} = 60.71\%$) significantly improves performance ($CV_{combo3} = 70.00\%$; $CV_{increase} = 9.28$; $P < 0.05$, Bonferroni corrected). Second, adding the right superior temporal pole to the combination of inferior frontal gyrus pars orbitalis (IF Orb) and inferior parietal cortex (IPar; $CV_{combo2} = 62.00\%$) significantly improved performance ($CV_{combo3} = 74.71\%$; $CV_{increase} = 12.71$; $P < 0.05$, Bonferroni corrected).

Combinatoric Analyses Identify Regions Contributing Unique Predictive Power. The ability to combine information from multiple regions is a fundamental strength of our multivariate approach. However, these analyses raise important questions concerning the contributions of each region: one region might consistently add information to other regions [unique combinatorial performance (UCP); Eq. S1], whereas another region might have the opposite

Smith et al.

www.manaraa.com

pattern, consistently gaining information when other regions are added to it [average combinatorial improvement (ACI); Eq. S2]. Importantly, this combinatoric approach provides a quantitative procedure for assessing both the UCP and ACI for each region after accounting for the influence of lesion size.

To account for lesion size and quantify the informational contributions of each of the 12 critical perisylvian regions that have previously been associated with spatial neglect in the literature (Table S1), we examined all combinations of two ROIs after lesion size was factored into the performance (Fig. 3A and Fig. S2). We found that the STG consistently improved the performance of other regions (mean increase = 4.94; $t_{(11)}$ = 4.35, $P < 0.02$ corrected), even after accounting for lesion size (Fig. 3 B and D, red). Importantly, examining all pairwise comparisons revealed that the right STG added more predictive power than several other regions, including the right supramarginal gyrus ($t_{(11)}$ = 3.81), the three divisions of the right inferior frontal gyrus (IF Orb: $t_{(11)}$ = 4.62; IF Tri: $t_{(11)}$ = 4.47; IF Oper: $t_{(11)}$ = 4.38), and the right inferior parietal lobe (IPar; $t_{(11)}$ = 3.92), indicating STG contains unique information for predicting neglect. We also found that the angular gyrus (mean increase = 4.33; $t_{(11)}$ = 4.27, $P < 0.02$ corrected) and insula (mean increase = 3.39; $t_{(11)}$ = 4.17, $P < 0.02$ corrected) consistently received predictive power from other

regions (Fig. 3 C and D, blue). Furthermore, examining all pairwise comparisons revealed that the angular gyrus was missing more information than the STG ($t_{(11)}$ = 3.89, $P < 0.05$, Bonferroni corrected for 66 comparisons), potentially suggesting that the classic association between right angular gyrus and spatial neglect (e.g., ref. 25) may reflect a byproduct of analytical tools that were heretofore incapable of quantifying interregional interactions.

Discussion

Applying MVPA to clinical questions is already producing insights into the subtle anatomical abnormalities that distinguish controls and neurological patients (35–38). Here, we applied MVPA to situations where there is profound structural damage (e.g., where there is a complete destruction of anatomical regions, rather than subtle thinning). An important requirement for this application was the development of rigorous methods for determining whether a classifier is performing above chance and determining whether additional predictors (e.g., ROIs and lesion volume) significantly improve classification performance. These considerations are critical, because simple measures alone—such as lesion volume or symptom incidence—often allow accurate classification levels for neurological or neuropsychological disorders far in excess of 50%. Our current results, constructed with the canonical example of spatial neglect after right hemisphere damage, showcase how this technique of simultaneously considering the influence of multiple regions or voxels can provide a unique approach for using lesion data to study human brain function.

Our combinatoric approach to the identification of diagnostic information yielded convergent findings for the right temporal cortex, particularly the STG. When we focused our selection of ROIs to those major right perisylvian cortical areas that have been observed to correlate with the core deficit of spatial neglect in various previous studies, namely the temporal-parietal junction within the inferior parietal lobe, the superior/middle temporal cortex and underlying insula, as well as the ventrolateral prefrontal cortex (for review, see ref. 8), the temporal region contributed the most information in predicting spatial neglect. This finding is enhanced by the complementary fact that the right frontal and/or parietal regions provided less predictive information. When combinations of 2 of the 12 ROIs were considered, we observed that adding the right STG to the orbital frontal cortex improved the prediction of spatial neglect compared with all regions in isolation. For combinations of three regions, we saw two cases where the initial two regions benefited from the addition of a third region. The combination of right inferior frontal and inferior parietal cortices was significantly improved by the addition of right middle or superior temporal pole areas. Moreover, when we examined the predictive power of each of the 12 critical right perisylvian regions when combined with one or two other ROIs after accounting for lesion size, we found that the STG consistently improved the classifier performance of other regions (Fig. 3). In other words, the right STG reliably contained unique information for predicting spatial neglect. In contrast, classifier performance for right inferior parietal and insular cortices was consistently enhanced by the addition of other ROIs, indicating these regions are lacking independent predictive information for spatial neglect. This latter finding suggests that these regions' involvement with spatial neglect could reflect an epiphenomenal finding in traditional univariate studies, as these traditional methods are incapable of quantifying interregional interactions.

Although our results demonstrate clear differences within regions in the right hemisphere, the present observations do not allow conclusions about a possible involvement of the human left hemisphere in spatial attention. However, we note that only a small percentage of patients with acute left hemisphere brain damage show spatial neglect (39, 40). Spatial orienting and
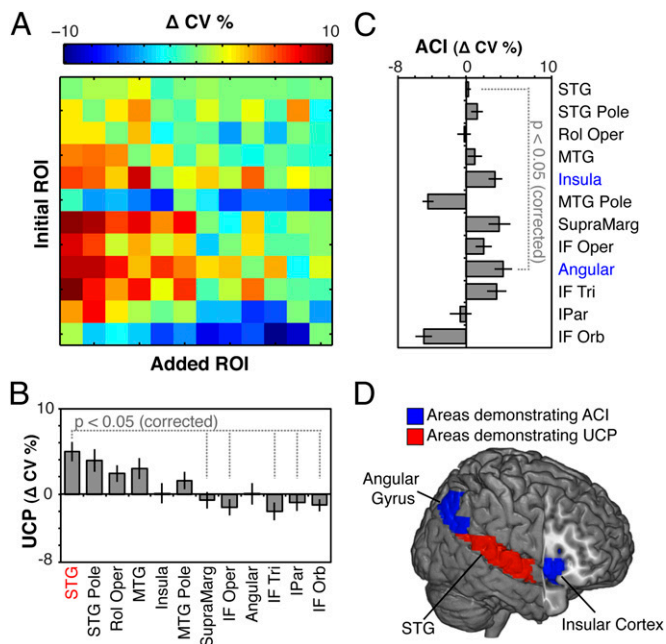

Fig. 3. Combinatoric analyses control for lesion size while revealing robustly predictive brain regions. (A) To examine the interaction of information contained in each of the 12 critical perisylvian ROIs in the right hemisphere that were previously associated with spatial neglect in the literature, we examined all combinations of two ROIs. We then computed, iteratively, the change in CV percentages when adding ROIs to each other. Crucially, these changes in CV percentages explicitly factor in lesion size, as this feature is present in all cells of the 12 × 12 matrix. (B) The STG was the only region that consistently added predictive power to other ROIs, indicating significant unique combinatorial performance (UCP) (Eq. S1). Pairwise comparisons revealed that UCP for STG was greater than several other regions, including the supramarginal gyrus, the three divisions of the right inferior frontal gyrus (pars opercularis: IF Oper; pars triangularis: IF Tri; pars orbitalis: IF Orb), and the inferior parietal lobe (IPar). (C) In contrast, classifier performance for two regions, angular gyrus and insula, was consistently improved by the addition of other ROIs, indicating significant average combinatorial improvement (ACI) (Eq. S2). Pairwise comparisons revealed that ACI for angular gyrus was greater than STG. (D) Summary of brain regions demonstrating significant UCP (red) and ACI (blue). Error bars in B and C indicate SEM.

attention thus appear to be dominantly represented in the human right hemisphere.

Our findings regarding the STG mesh with known structural and functional architecture. The right temporal cortex appears to play a key role in integrating a densely interconnected perisylvian network representing spatial orienting and exploration in humans. This region may function as a central hub: connecting with the inferior parietal lobe via posterior parts of the middle longitudinal fasciculus (MdLF) and extreme capsule/inferior occipitofrontal fasciculus (EmC/IOF), as well as connecting with the lateral prefrontal cortex via the arcuate fasciculus (AF) and the EmC/IOF (for review, see ref. 41). Beyond its connection with the dorsal route of visual information processing, the superior temporal gyrus also receives polysensory inputs from the ventral perceptual stream. When different tracers were injected into the posterior parietal and inferotemporal cortex of the same hemisphere, overlapping labeling was found near the fundus of the superior temporal sulcus (42, 43). The superior temporal cortex thus represents a site for multimodal sensory convergence (for review, see ref. 44), with converging inputs from the dorsal and ventral streams.

Regional analysis focused on those 12 cortical right perisylvian ROIs identified previously by numerous studies, techniques, and groups as correlating with spatial neglect (for review, see refs. 8 and 9). However, our analysis also included a more comprehensive set of 45 regions facilitating a richer perspective that considers all usable features (i.e., voxels that are not damaged or spared across all patients). Considering all usable features is important for two reasons. First, it allowed us to identify the influence of regions on spatial neglect without any anatomical a priori assumptions. Second, this approach combined both positive predictors (regions where injury predicts a deficit) with negative predictors (where injury predicts spared abilities), potentially providing a more accurate classification than an analysis that is restricted to positive predictors (7). We found a combination of three regions provided a robust predictor (>86%) for spatial neglect. Specifically, we found that the amount of injury to the right middle occipital cortex, Heschl's gyrus in the superior temporal cortex, and the corticospinal tract—while controlling for overall lesion size—were able to reliably predict neglect better than any univariate measure and better than any region combined with lesion volume. This is a surprising result, as each of these regions is strongly associated with a function other than spatial attention. There are several candidate reasons for this apparently counterintuitive result.

Consider the following explanations. Whereas we found the best predictors appear outside the regions classically associated with the spatial neglect, our findings also reveal that multivariate analyses can provide accurate classification based solely on the size and extent of brain injury, as observed on standard clinical scans. One parsimonious explanation for this apparent paradox is that the multivariate analyses are able to capitalize on the regions that are both positively predictive and negatively predictive of the disorder (7). The best independent classifiers may largely provide redundant information with lesion volume. Conversely, some of the best classifiers for analyses that include additional predictors such as overall lesion volume may be ones that predict the absence of a symptom when they are injured while overall lesion volume is small. Indeed, the middle occipital cortex and corticospinal tract show this pattern, as they lie along the border of the right middle cerebral artery territory that has classically been associated with spatial neglect.

A further benefit of studying all usable features is the ability to harness both positive and negative predictors, because negative predictors are likely to be outside the standard set of regions discussed. This dual predictive power is likely to improve classification: brain injury data are noisy (functional modules are typically larger than voxels, multiple nodes can influence

behavior, precise location of function can vary across individuals), so these properties should lead to better classification and generalizability than is possible with mass univariate approaches. Of course, generalizability of a predictive model is always limited by the within-sample variety of the training data, but our present sample includes a wide range of damage across the right hemisphere (Fig. 1). A possible application of such a prediction in a clinical context might be an improved prediction of long-term severity of the disorder under study, such as recovery and rehabilitation processes, based on the structural scans taken in the acute stage of the stroke (for an example of such a prediction strategy using the traditional VLSM approach, see ref. 45). Future work that combines the methods described here (acute scans with acute behavioral measures) with chronic behavioral measures may prove to be an important tool for long-term prognosis. Such studies could triage individuals who will spontaneously recover from those who will need focused rehabilitation and those who are unlikely to recover function regardless of treatment. Effectively, this could help differentiate between brain regions that are not involved with a given task, those that are critical but with redundancy, and those where there are no mechanisms for compensation.

To conclude, the method we outline has broad implications for cognitive neuroscience, providing a unique approach for using lesion data to study human brain function. MVPA in lesion mapping will be a useful tool for the wide spectrum of cognitive functions and clinical disorders where—as in spatial neglect as a motivating example—similar symptoms can be observed following discrete injury to portions of a network. Our approach could also be extended using existing algorithms, including regression-based machine learning (46), to examine continuous measures, such as neglect severity (47). Additionally, as learning algorithms and computing efficiency improve, the reliance on arbitrary ROI definitions (a limitation of our approach) can be abandoned by fractionating the brain into specific segments of varying size and spatial contiguity; this extension could provide an opportunity to gain even greater insight into small- and large-scale network-driven behavior.

Materials and Methods

Patient Sample. Data were drawn from 140 consecutively admitted stroke patients with circumscribed right hemisphere lesions (*SI Materials and Methods*) from a well-defined recruitment area belonging to the University of Tuebingen, sampled over a period of 7 y. All patients gave their informed consent to participate in the study, which has been performed in accordance with the ethical standards in the 1964 Declaration of Helsinki and approved by the institutional review board of the University of Tübingen. The full details regarding patient characteristics, test criteria, and other demographic information are reported in our previous work, which uses the same sample of patients (20). In short, following standardized testing for spatial neglect by using the letter cancellation task, the Bells test, the baking tray task, and a copying task, the subjects were divided into a group of 78 patients with spatial neglect and 62 control patients who did not show the disorder. Neglect patients fulfilled the criterion for spatial neglect in at least two of these four clinical tests and showed the typical spontaneous deviation of eyes and head toward the ipsilesional side and ignoring of contralesionally located people or objects (20).

Feature Space Construction. For effective machine learning, the analytic procedure behind MVPA (Fig. 1A) geared toward prediction, the classifier (learning algorithm) generally uses multiple pieces of information, or features, to use to identify similarities within or differences between the classes that are being predicted (34). We created a 4D dataset (a 3D image of each individual patient across the sample, n = 140) of 91 × 109 × 91 voxels (downsampled by a factor of 2 from the original resolution to reduce overfitting). The union of all patient lesions (Fig. 1B, green) consisted of 20,068 voxels (2-mm³ resolution), whereas the set of voxels that were lesioned in at least 5% of the patient population was 12,811 voxels (2-mm³ resolution). None of these voxels were lesioned across all patients, an observation that would preclude their use in our analyses. To reduce the likelihood of the classifier capitalizing on rare events, our primary analyses

www.manaraa.com

focused on voxels that were damaged in at least 5% of patients. Then, using regions that were damaged in at least 5% of patients, we also constructed more selective feature spaces ($n = 45$) using brain ROIs from the AAL (highlights gray matter, see ref. 48) and Juelich (white matter, see ref. 49) atlases. These brain regions vary in size, as they are defined anatomically (Table S1). We note that 11 regions (from our original 56 in Fig. 1C) were excluded because no voxel within them was damaged in at least 5% of patients (Table S1). We also examined all combinations two and three ROIs (SI Materials and Methods).

Classification. Classification analyses of the structural images were completed using PyMVPA (50). We used two classifiers to develop models for predicting neglect: a linear SVM ($\nu = 0.5$; C was scaled between 0 and 1, depending on the distribution of data) and a nonlinear [radial basis function kernel (RBF)] SVM ($\nu = 0.5$; C = 1; $\gamma = 0.5$). Although there was a scant amount of re-dundancy and correlation between our features (Fig. S4), we note that, unlike some classifiers (e.g., naïve Bayes), SVM classifiers do not assume in-dependence between features (34, 51). Performing classifications with non-independent features is unavoidable for lesion-mapping approaches, because larger lesions could be associated with correlated patterns of damage in specific areas. Additionally, our soft-margin SVM relies on maximizing the separation between the margin and decision boundary, and using a relatively small C should help limit susceptibility to overfitting in high-dimensional problems (34), an important consideration given our number of features and observations. To avoid constructing a biased model, all training sets consisted of equal observations of neglect and control patients. As this procedure entailed dropping a random subset of observations from each classification, we repeated each classification five times to ensure a stable CV without any upward bias (52). We used the leave-one-out CV percentage as our metric for evaluating classifier performance. The constructed algorithm and its set of weights were then used to make a prediction on the left-out individual. It-eratively, the leave-one-out process was applied to each individual. The CV percentage is then the average of the five iterations of all 140 predictions. Significance of CVs and changes in CV were assessed using Monte Carlo permutation-based statistical testing (SI Materials and Methods) (53).

1. Fellows LK, et al. (2005) Method matters: An empirical study of impact in cognitive neuroscience. *J Cogn Neurosci* 17(6):850–858.
2. Rorden C, Karnath H-O (2004) Using human brain lesions to infer function: A relic from a past era in the fMRI age? *Nat Rev Neurosci* 5(10):813–819.
3. Bates E, et al. (2003) Voxel-based lesion-symptom mapping. *Nat Neurosci* 6(5):448–450.
4. Rorden C, Fridriksson J, Karnath HO (2009) An evaluation of traditional and novel tools for lesion behavior mapping. *Neuroimage* 44(4):1355–1362.
5. Haynes JD, Rees G (2006) Decoding mental states from brain activity in humans. *Nat Rev Neurosci* 7(7):523–534.
6. Norman KA, Polyn SM, Detre GJ, Haxby JV (2006) Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10(9):424–430.
7. Kriegeskorte N, Bandettini P (2007) Analyzing for information, not activation, to exploit high-resolution fMRI. *Neuroimage* 38(4):649–662.
8. Karnath HO, Rorden C (2012) The anatomy of spatial neglect. *Neuropsychologia* 50(6):1010–1017.
9. Corbetta M, Shulman GL (2011) Spatial neglect and attention networks. *Annu Rev Neurosci* 34(1):569–599.
10. Karnath HO, Dieterich M (2006) Spatial neglect—A vestibular disorder? *Brain* 129(Pt 2):293–305.
11. Becker E, Karnath H-O (2010) Neuroimaging of eye position reveals spatial neglect. *Brain* 133(Pt 3):909–914.
12. Heilman KM, Bowers D, Watson RT (1983) Performance on hemispatial pointing task by patients with neglect syndrome. *Neurology* 33(5):661–664.
13. Behrmann M, Watt S, Black SE, Barton JJS (1997) Impaired visual search in patients with unilateral neglect: an oculographic analysis. *Neuropsychologia* 35(11):1445–1458.
14. Karnath H-O, Niemeier M, Dichgans J (1998) Space exploration in neglect. *Brain* 121(Pt 12):2357–2367.
15. Jelsone-Swain LM, Smith DV, Baylis GC (2012) The effect of stimulus duration and motor response in hemispatial neglect during a visual search task. *PLoS ONE* 7(5):e37369.
16. Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven atten-tion in the brain. *Nat Rev Neurosci* 3(3):201–215.
17. Thiebaut de Schotten M, et al. (2011) A lateralized brain network for visuospatial attention. *Nat Neurosci* 14(10):1245–1246.
18. Smith DV, et al. (2010) Spatial attention evokes similar activation patterns for visual and auditory stimuli. *J Cogn Neurosci* 22(2):347–361.
19. Karnath HO, Ferber S, Himmelbach M (2001) Spatial awareness is a function of the temporal not the posterior parietal lobe. *Nature* 411(6840):950–953.
20. Karnath H-O, Fruhmann Berger M, Küker W, Rorden C (2004) The anatomy of spatial neglect based on voxelwise statistical analysis: A study of 140 patients. *Cereb Cortex* 14(10):1164–1172.
21. Corbetta M, Kincade MJ, Lewis C, Snyder AZ, Sapir A (2005) Neural basis and recovery of spatial attention deficits in spatial neglect. *Nat Neurosci* 8(11):1603–1610.
22. Chechlacz M, et al. (2010) Separating neural correlates of allocentric and egocentric neglect: Distinct cortical sites and common white matter disconnections. *Cogn Neuropsychol* 27(3):277–303.
23. Rengachary J, He BJ, Shulman G, Corbetta M (2011) A behavioral analysis of spatial neglect and its recovery after stroke. *Frontiers Hum Neurosci* 5:29.
24. Heilman KM, Watson RT, Valenstein E, Damasio AR (1983) Localization of lesions in neglect. *Localization in Neuropsychology*, ed Kertesz A (Academic Press, New York), pp 471–492.
25. Mort DJ, et al. (2003) The anatomy of visual neglect. *Brain* 126(Pt 9):1986–1997.
26. Husain M, Kennard C (1996) Visual neglect associated with frontal lobe infarction. *J Neurol* 243(9):652–657.
27. He BJ, et al. (2007) Breakdown of functional connectivity in frontoparietal networks underlies behavioral deficits in spatial neglect. *Neuron* 53(6):905–918.
28. Karnath HO, Rorden C, Ticini LF (2009) Damage to white matter fiber tracts in acute spatial neglect. *Cereb Cortex* 19(10):2331–2337.
29. Urbanski M, et al. (2008) Brain networks of spatial awareness: Evidence from diffusion tensor imaging tractography. *J Neurol Neurosurg Psychiatry* 79(5):598–601.
30. Shinoura N, et al. (2009) Damage to the right superior longitudinal fasciculus in the inferior parietal lobe plays a role in spatial neglect. *Neuropsychologia* 47(12):2600–2603.
31. Clithero JA, Carter RM, Huettel SA (2009) Local pattern classification differentiates processes of economic valuation. *Neuroimage* 45(4):1329–1338.
32. Hampton AN, O'Doherty JP (2007) Decoding the neural substrates of reward-related decision making with functional MRI. *Proc Natl Acad Sci USA* 104(4):1377–1382.
33. Carter RM, Bowling DL, Reeck C, Huettel SA (2012) A distinct role of the temporal-parietal junction in predicting socially guided decisions. *Science* 337(6090):109–111.
34. Hastie T, Tibshirani R, Friedman J (2009) *The Elements of Statistical Learning* (Springer, New York).
35. Koutsouleris N, et al. (2009) Use of neuroanatomical pattern classification to identify subjects in at-risk mental states of psychosis and predict disease transition. *Arch Gen Psychiatry* 66(7):700–712.
36. Gothelf D, et al. (2011) Developmental changes in multivariate neuroanatomical patterns that predict risk for psychosis in 22q11.2 deletion syndrome. *J Psychiatr Res* 45(3):322–331.
37. Sun D, et al. (2009) Elucidating a magnetic resonance imaging-based neuroanatomic biomarker for psychosis: Classification analysis using probabilistic brain atlas and machine learning algorithms. *Biol Psychiatry* 66(11):1055–1060.
38. Hackmack K, Paul F, Weygandt M, Allefeld C, Haynes JD; Alzheimer's Disease Neu-roimaging Initiative (2012) Multi-scale classification of disease using structural MRI and wavelet transform. *Neuroimage* 62(1):48–58.
39. Suchan J, Karnath HO (2011) Spatial orienting by left hemisphere language areas: a relict from the past? *Brain* 134(Pt 10):3059–3070.
40. Becker E, Karnath HO (2007) Incidence of visual extinction after left versus right hemisphere stroke. *Stroke* 38(12):3172–3174.
41. Karnath HO (2009) A right perisylvian neural network for human spatial orienting. *The Cognitive Neurosciences*, ed Gazzaniga MS (MIT Press, Cambridge, MA), Vol IV, pp 259–268.
42. Morel A, Bullier J (1990) Anatomical segregation of two cortical visual pathways in the macaque monkey. *Vis Neurosci* 4(6):555–578.
43. Baizer JS, Ungerleider LG, Desimone R (1991) Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques. *J Neurosci* 11(1):168–190.
44. Karnath HO (2001) New insights into the functions of the superior temporal cortex. *Nat Rev Neurosci* 2(8):568–576.
45. Karnath H-O, Rennig J, Johannsen L, Rorden C (2011) The anatomy underlying acute versus chronic spatial neglect: A longitudinal study. *Brain* 134(Pt 3):903–912.
46. Cohen JR, et al. (2011) Decoding continuous behavioral variables from neuroimaging data. *Frontiers Neurosci* 5:75.
47. Rorden C, Karnath H-O (2010) A simple measure of neglect severity. *Neuro-psychologia* 48(9):2758–2763.
48. Tzourio-Mazoyer N, et al. (2002) Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage* 15(1):273–289.
49. Bürgel U, et al. (2006) White matter fiber tracts of the human brain: Three-di-mensional mapping at microscopic resolution, topography and intersubject variabil-ity. *Neuroimage* 29(4):1092–1105.
50. Hanke M, et al. (2009) PyMVPA: A unifying approach to the analysis of neuroscientific data. *Front Neuroinform* 3:3.
51. Schölkopf B, Smola AJ (2001) *Learning with Kernels: Support Vector Machines, Reg-ularization, Optimization, and Beyond* (MIT Press, Cambridge, MA).
52. Kim J-H (2009) Estimating classification error rate: Repeated cross-validation, re-peated hold-out and bootstrap. *Comput Stat Data Anal* 53(11):3735–3745.
53. Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neu-roimaging: A primer with examples. *Hum Brain Mapp* 15(1):1–25.

NEUROSCIENCE