Decoding the Brain's Surface to Track Deeper Activity

Mark L. Tenzer [1], Jonathan M. Lisinski [1] and Stephen M. LaConte [1,2]*

[1] Fralin Biomedical Research Institute at VTC, Virginia Tech, Roanoke, VA, United States, [2] Department of Biomedical Engineering and Mechanics, Virginia Tech, Blacksburg, VA, United States

Neural activity can be readily and non-invasively recorded from the scalp using electromagnetic and optical signals, but unfortunately all scalp-based techniques have depth-dependent sensitivities. We hypothesize, though, that the cortex's connectivity with the rest of the brain could serve to construct proxy signals of deeper brain activity. For example, functional magnetic resonance imaging (fMRI)-derived models that link surface connectivity to deeper regions could subsequently extend the depth capabilities of other modalities. Thus, as a first step toward this goal, this study examines whether or not surface-limited support vector regression of resting-state fMRI can indeed track deeper regions and distributed networks in independent data. Our results demonstrate that depth-limited fMRI signals can in fact be calibrated to report ongoing activity of deeper brain structures. Although much future work remains to be done, the present study suggests that scalp recordings have the potential to ultimately overcome their intrinsic physical limitations by utilizing the multivariate information exchanged between the surface and the rest of the brain.

Keywords: support vector machine, resting state connectivity, functional magnetic resonance imaging, multimodal, cerebral cortex

1. INTRODUCTION

Every neuroimaging method has intrinsic physical limitations. But even once the technology is sufficiently advanced to harness the theoretical physical capabilities of a given approach, there is often room for sometimes astounding additional methodological and algorithmic innovation. As an example, today's functional magnetic resonance imaging (fMRI) has benefited from major transformations that surpassed the perceived physical limits in terms of both resolution and acquisition speed. Specifically, MRI exceeds the diffraction limits of 43 MHz/T resolution (an approximate field-dependent wavelength of 2 m·T) by multiple orders of magnitude thanks to the insights of Paul Lauterbur that magnetic field gradients can be used to spatially encode the NMR signal and, in addition, that algorithmic approaches can be used to reconstruct images from such data (Lauterbur, 1973). The use of magnetic gradients, though, makes MRI slow. Nonetheless major speed gains have been realized over the years by repeatedly looking beyond the physical limits of the current conventions. Echoplanar imaging (Mansfield, 1977), in-plane parallel imaging (Sodickson and Manning, 1997; Pruessmann et al., 1999; Griswold et al., 2002), and simultaneous multi-slice acquisitions (Larkman et al., 2001; Nunes et al., 2006; Feinberg et al., 2010; Moeller et al., 2010) have all culminated to exceed previous "speed limits" by orders of magnitude and to enable sub-second sampling rates.

Although fMRI is a unique and particularly successful example, one general approach for circumventing physical limitations has been to look to multimodal imaging to merge two or more techniques with complementary capabilities. In human studies, common modalities for functional measurements include magnetoencephalography (MEG), electroencephalography (EEG), positron emission tomography (PET), functional near-infrared spectroscopy (fNIRS), and fMRI. Multimodal examples include combining structural and functional imaging (e.g., using anatomical MRI to constrain source localization in EEG Pascual-Marqui, 1999; Fiederer et al., 2016), as well as combining shared primary physiological measures (e.g., obtaining both spatial and temporal resolution from fMRI and EEG or MEG (Menon et al., 1997; Dale et al., 2000). The most common motivation for employing multimodal techniques is to simultaneously capture the highest possible level of both temporal and spatial resolution (Meyer-Lindenberg, 2010). *Depth of measurement*, however, is a related physical limitation of scalp-based modalities that has been largely overlooked in the multimodal literature.

To simulate approaches to bypass physical limitations in measurement depth, we used resting-state fMRI data to examine whether surface activity could serve as a proxy for deeper brain signals. The crux of this idea is the hypothesis that the multivariate connectivity of the brain "transmits" enough information to the surface of cortex to enable decoding of target signals from interior regions and networks. This study tests that basic premise by examining whether whole-brain measurements could be used to calibrate surface-restricted ones. The general concept of training a surface-limited model to track a deeper target signal is illustrated in **Figure 1**. Namely, a training step produces a model comprised of multivariate weight vectors that can subsequently track the target signal using only recordings from the surface. Arguably, this fMRI-based initial demonstration of feasibility would be most easily translatable to a hemodynamic measure, like fNIRS. Nevertheless, future studies will be needed to evaluate strategies for fMRI to calibrate surface

modalities like EEG, fNIRS, and MEG. To fully implement a multimodal strategy around this idea will ultimately require multiple steps. Perhaps the most challenging of which would be to evaluate forward and reverse transfer functions between fMRI and the other modalities (e.g., EEG's correspondence between hemodynamic responses; Sato et al., 2010).

Currently, it is not known if even surface fMRI signals can reconstruct deeper fMRI signals. Thus, this study aimed to demonstrate that such an approach can accurately track such localized and distributed signals. Our primary question was "Can we use the surface of the brain to track distributed networks and anatomical regions?" This was examined through cross-validation on a thirteen-participant data set under two conditions i) a fixed-depth surface mask, and ii) a sparse, low resolution mask of strategic surface locations that was specific for each of the 16 target signals. The sparse masks were generated with independent data from 99 participants, which additionally enabled us to explore our secondary question "Can we map the surface connectivity that enables decoding?" Finally, we examined the reproducibility of our results by repeating our cross-validation estimates in 83 participants from an archival data set (openneuro.org; Power et al., 2017). Together, our findings suggest that surface techniques could enjoy a fundamental advancement by overcoming their intrinsic physical limitations through the use of machine learning models that capture whole brain functional connectivity.

2. MATERIALS AND METHODS

Participants and Data Collection

A group of 13 participants (eight females, mean age 26 years) was used to estimate cross-validated prediction accuracies. A second, larger group of 99 participants (48 females; mean age 25 years) was formed from six ongoing protocols in the lab and used to construct group-level support vector regression (SVR) maps and volumetrically-subdivided feature masks. All participants gave

FIGURE 1 | Illustration of the surface decoding approach. Left: At each point in time, voxels from the surface of the brain (yellow) are extracted. Their intensities comprise the training features, $\vec{x}_{train}(t)$. In addition, training labels, $y_{train}(t)$, are extracted from the average of voxels in a network or region of interest (green). (Not illustrated: Induction is used to train a regression model, $\vec{w}$, that relates $\vec{x}_{train}(t)$ to $y_{train}(t)$). Right: In a new dataset, $\vec{x}_{test}(t)$ and $y_{test}(t)$ are extracted. Combining each testing vector with the regression weights $\vec{w}$ yields the predicted testing label $\hat{y}_{test}(t)$ (brown). In this study, the correlation between the predictions, $\hat{y}_{test}(t)$, and the true values, $y_{test}(t)$, is defined as the prediction accuracy.

informed consent, and the study was done in accordance with the Institutional Review Board of Virginia Tech.

Structural and functional brain data were acquired on a 3 Tesla whole-body scanner (Siemens Trio). A three-dimensional magnetization prepared rapid acquisition gradient echo (MPRAGE) pulse sequence (Mugler and Brookeman, 1990) was used to collect T1-weighted anatomical volumes (resolution = $1 \times 1 \times 1$ mm^3; TR = 2,600 ms; TE = 3.02 ms; TI = 900 ms, FOV = 256 mm^2; FA = 8°). Echo planar imaging (EPI) was used to acquire resting-state data. For the 13-participant data set, two 6 min runs were acquired (33 interleaved axial slices; resolution = $3.4 \times 3.4 \times 3.6$ mm^3; TR = 2,000 ms; TE = 30 ms; FOV = 220 mm^2; FA = 90°). For the 99 participants, only one resting state run was used per participant and there was minor variability across acquisition protocols with slice thickness, number of slices, TR, and scan duration ranging between 3.6 and 4 mm, 29 and 33 interleaved axial slices, 1,750 and 2,000 ms, and 6–10 min., respectively. Specifically: $n = 56$, slice thickness = 3.6 mm, 33 slices, TR = 2,000 ms, duration = 6.1 min; $n = 36$, slice thickness = 3.6 mm, 29 slices, TR = 1,850 ms, duration = 10.0 min; $n = 7$, slice thickness = 4.0 mm, 33 slices, TR = 2,000 ms, duration = 6.0 min. The resting state instructions were consistent throughout. Participants were instructed to keep their eyes open during the scan and direct their gaze on a white plus (+) sign centered on a black background. The verbal instructions before the scan were, "Remember to keep your eyes open and directed at the fixation symbol. Let your mind wander freely. If you notice yourself focusing on any particular train of thought, let your mind wander away from it." Note that the **Supplementary Material** characterizes the quality of these two datasets. **Supplementary Figures S1–S4** provide an overview of motion in terms of both displacement and displacement change for the $N = 99$ and $N = 13$ datasets.

Finally, while our findings were statistically robust, our 13-subject cohort is admittedly small. To evaluate the reproducibility of our findings we used 83 participants from an archival dataset (openneuro.org; Power et al., 2017) that included two resting-state runs per participant to replicate our cross-validated prediction accuracy estimates. These data were acquired at 3T (Siemens MAGNETOM Trio) with TR = 2,500 ms; TE = 27 ms; voxels = 4 x 4 x 4 mm; 32 interleaved slices. Run lengths varied from 130 to 133 volumes (mean: 132 volumes).

Surface Masks

As shown in **Figure 1** the training features are depth-limited signals. As such, this approach could be considered a surface-based analysis. The goal however is not to use sophisticated surface techniques such as those supported by FreeSurfer (Fischl, 2012) to reconstruct the topology of the cortical surface. Rather, the focus here for evaluating the potential future application to scalp-based measurements simply requires depth-limited masks of the brain's folded surface to simulate depth-attenuated measurements.

We used morphological image processing on binary whole-brain masks to obtain surface cortical masks. These surface masks excluded regions that would not be easily accessible to surface-based measurements (e.g., the inferior surface of the brain). The approach for deriving the surface masks is illustrated in **Figure 2**. Surface masks were developed from whole-brain binary masks that were zero-padded in the superior and inferior planes (Mask A). For mask thickness m, a "subsurface" mask was generated by zeroing each voxel within m millimeters from the surface boundary in any direction. That is, a new binary mask was produced that excluded voxels outside the brain as well as within m millimeters of the surface of the brain (see Mask B). Next a brain surface mask was generated by including all voxels within the whole-brain mask while excluding all voxels within the subsurface mask (Mask A − Mask B).

Finally, to exclude surface regions that are not easily accessible to surface-based methods (e.g., the cerebellum and the inferior surface of the brain), the subsurface mask was repeatedly translated in the inferior direction, and after the first shift, each

FIGURE 2 | Masking the surface. First, a whole-brain mask was generated (A). All voxels within mask thickness m of the surface were eliminated to generate (B), containing the interior of the brain. A surface mask was then generated (A − B). To eliminate the inferior surface, (B) was repeatedly translated downward along the z-axis. After the first shift, each subsequent translational iteration was added to the previous one, creating an extended inferior volume shown in (C). Finally, C was used to exclude inferior regions in the final mask (A − B − C).

subsequent translational iteration was added to the previous one, creating an extended inferior volume (Mask *C*). This effectively removed all voxels inferior to the cortex (Mask *A* − Mask *B* − Mask *C*). Note that for this study, we fixed the surface mask thickness, *m*, to 10 mm, which represents a conservative depth that is applicable to both MEG and fNIRS to explore the utility of surface-based tracking.

Preprocessing

We used tools from AFNI (Cox, 1996) and FSL (Smith et al., 2004) to preprocess the imaging data. Specifically, FLIRT (Jenkinson et al., 2002) and FNIRT were used to register skull-stripped anatomical images to the MNI 152 space. Skull stripping was performed in AFNI and visually inspected for errors. We generated white matter and cerebrospinal fluid masks using a probability ≥ 0.99 threshold to both corresponding maps generated by the hidden Markov random field model methods implemented in FAST (Zhang et al., 2001). Resting-state data were slice-timing and motion corrected using the 3dTshift and 3dvolreg commands in AFNI, respectively. We used the mean image from each run to derive a linear transformation matrix to co-register each functional scan to that participant's corresponding anatomic image when using FLIRT. Next, nuisance variable regression was performed by regressing out white matter and cerebrospinal fluid time series, the six motion parameters calculated during motion correction, and a 4th order polynomial to account for baseline drifts (Friston et al., 1996; Lund et al., 2006). The functional-to-anatomical and anatomical-to-MNI152 transformations calculated for each dataset were concatenated to construct a functional-to-MNI152 transform. This transform was applied to the functional images, producing 4 mm isotropic data in MNI space. These images were then smoothed using a 6 mm^2 full-width-at-half-maximum (FWHM) Gaussian kernel.

Support Vector Regression

SVR models for the masked resting-state datasets were generated using 3dsvm (LaConte et al., 2005), a wrapper of SVM^{light} (Joachims, 1999) that is convenient for processing fMRI data. Our general approach was to normalize each subject's T1 anatomical volume to MNI space and extract a target time series from either a network or anatomical region. The target time series were generated from both resting state networks (RSNs) and atlas-defined regions, and the features consisted of either brain surface voxels or surface group map-based volumetric subdivisions. For the RSNs, we used the top 10 components from a 20-component ICA meta-analysis of resting state as well as task data by Smith et al. (2009) and obtained the target time series using spatial regression. The atlas-defined bilateral regions were extracted from the Eickhoff et al. (2005) macro-label atlas available in AFNI. We examined amygdala, anterior cingulate cortex (ACC), caudate nucleus, insula, posterior cingulate cortex (PCC), and putamen. The time series for the atlas-defined regions were calculated by extracting the ROI time series data from MNI space at 4 mm^3 resolution and by averaging the time series data within each ROI. The time series targets served as labels for the

surface features to generate SVR models using 3dsvm (Joachims, 1999; LaConte et al., 2005).

Estimating Prediction Accuracy

Using the 13-participant dataset, each participant's two resting-state runs alternately served as training and testing data to produce a cross-validated (averaged) estimate of prediction accuracy. We defined prediction accuracy as the Pearson's correlation coefficient *r* between predicted and observed target time series. Appropriate SVR parameters were evaluated by Craddock et al. (2013); based on our previous work, *C* was set to 100 and *ε* was set to 0.1 with a linear kernel for surface-masked data.

To test the statistical significance of the prediction accuracies, we used a "wavestrapping" algorithm as outlined by previous groups (Breakspear et al., 2004; Bullmore et al., 2004). Specifically, a discrete wavelet decomposition of true target time series was used to generate 26,000 surrogate time series, that in turn were used to train SVR models and then predict the 'true' labels in the appropriate test data. The resulting set of SVR-based correlations provided a null distribution to non-parametrically estimate the *p*-value for each network and bilateral region. In other words, the "true training label"-based model was applied to the test data to produce a predicted time series of correlation accuracy *r*. Then the significance of *r* was estimated by training with the distribution of wavestrapped training labels to produce a distribution of correlations with the (unaltered) test data labels. Specifically the non-parametric *p*-value was approximated as the proportion of trials in which the true target's correlation exceeded the distribution of surrogate correlations.

The PyWavelets Python module (Wasilewski, 2006) was used to perform discrete wavelet decompositions on the target time series using the 4th-order Daubechies wavelet. Four levels were used, based on the equation

$$max\_level = \left\lfloor \log_2 \frac{\text{signal length}}{\text{filter length} - 1} \right\rfloor, \qquad (1)$$

which ensures the signal is longer than the FIR filter length of the wavelet. Here, the signal contained 182 measurements and the decomposition filter length of the 4th-order Daubechies wavelet is 8. The detail coefficients, but not the approximation coefficients, at each level were randomly permuted. Then the resulting surrogate target time series was reconstructed with an inverse wavelet transform. As mentioned, the null distribution was constructed by repeating this process for 26,000 permutations for every network and region (an average of 1,000 permutations per run for every network and region) using GNU parallel (Tange, 2011). Finally, to evaluate the reproducibility of these results, we also estimated *r* values using archival data from 83 participants (openneuro.org; Power et al., 2017).

Group Maps and Target-Specific Surface Subdivisions

In the 99-participant data set, a single resting-state run was used to generate surface SVR models. These results were subsequently used as first-level data to generate group maps for every RSN and

bilateral region. The significant clusters within these group maps were then volumetrically subdivided to derive a sparse feature set for each target RSN and region. Each group map was organized into clusters of five or more nonzero voxels. Each voxel had to be within one grid cell of its nonzero neighbor to be included in the cluster. These clusters were recursively subdivided to ensure that their total volumes would be less than or equal to a fixed maximum volume, set to 40 voxels ($2,560 \, \text{mm}^3$). To do this, the two points within a cluster with the greatest distance between them were found; these were used to derive a vector between them, $\boldsymbol{v}$, and a midpoint, $\boldsymbol{x}$. These were used to construct a plane normal to $\boldsymbol{v}$ and containing $\boldsymbol{x}$, bisecting the cluster: $\boldsymbol{v} \cdot \boldsymbol{x} = b$. This procedure was used to generate two new clusters from each existing cluster. This subdivision procedure was repeated recursively through all clusters greater than 40 voxels in size. To use the subdivided surface as features, the voxel values within each subdivision were averaged. SVM^{light} was used directly and because these models had far fewer features, C and ϵ were set to the SVM^{light} defaults; C varied from model to model, whereas ϵ remained at 0.1.

Summary of Methods

To summarize, we used 10 RSNs from Smith et al. (2009) as well as six bilateral anatomical regions. SVR, restricted to a 10-mm-thick surface of the brain, was used to model these 16 target signals. Data from 13 participants who completed two resting-state runs provided the ability to estimate prediction accuracy using split-half resampling (training on one run and testing on the other) (Strother et al., 2002). This was then replicated with additional 83-participant archival data set. To generate group maps of the surface, SVR training was performed on an additional 99 independent participants who completed only one resting-state run. For each target, the 99 SVR maps were combined using a one-sample t-test. Finally, to examine the impact of sparser and lower spatial resolution measures, these maps were volumetrically subdivided to create feature sets to also test with the 13- and 83-participant data sets.

3. RESULTS

Depth of Measurement

This study used a mask thickness of 10 mm throughout. To characterize the impact of depth, however, **Figure 3** shows the 13 participant SVR accuracies averaged across the 10 RSNs and 6 regions for mask thicknesses over a range of 5 to 20 mm. As shown, prediction accuracy is stable and high over this range, with no peak optimum or sign of asymptotic plateau. We observed that accuracy increases monotonically, with correlations increasing by approximately 0.005 per mm of mask thickness ($R^2 = 0.953$). Interestingly, participants tend to rank by prediction accuracy. In other words, with only minor exceptions, the relative performance of each participant was consistent across all map thicknesses. Thus, some participants' brain activity was consistently easier to predict than others. This observation is examined further in our Supplemental Results (**Supplementary Figure S5**).

Sparsity of Surface Features

To examine sparser and lower spatial resolution measures, we volumetrically subdivided the SVR group maps for each of the 16 target signals. To do this, spatially distinct clusters were recursively subdivided to produce volumes of $2,560 \, \text{mm}^3$ (40 voxels) or less. The end result produced subdivided surface masks that were specific for each of the target signals. **Figure 4** shows an example of the recursive volumetric subdivision for a surface cluster spanning right inferior frontal, superior temporal, and inferior parietal regions (from the insula target SVR group map). The **Supplementary Tables S1–S16** include the volume and MNI coordinate for the centroid of each subdivision for all 16 target signals. For comparison, the full 10 mm surface models used approximately 8,000 voxels, and the subdivision produced an approximately 30–800 fold reduction of features (the number of subdivisions/reduced features were RSN1: 73, RSN2: 104, RSN3: 159, RSN4: 175, RSN5: 22, RSN6: 178, RSN7: 211, RSN8: 115, RSN9: 273, RSN10: 242, Amygdala: 10, ACC: 16, Caudate: 14, Insula: 50, PCC: 39, Putamen:
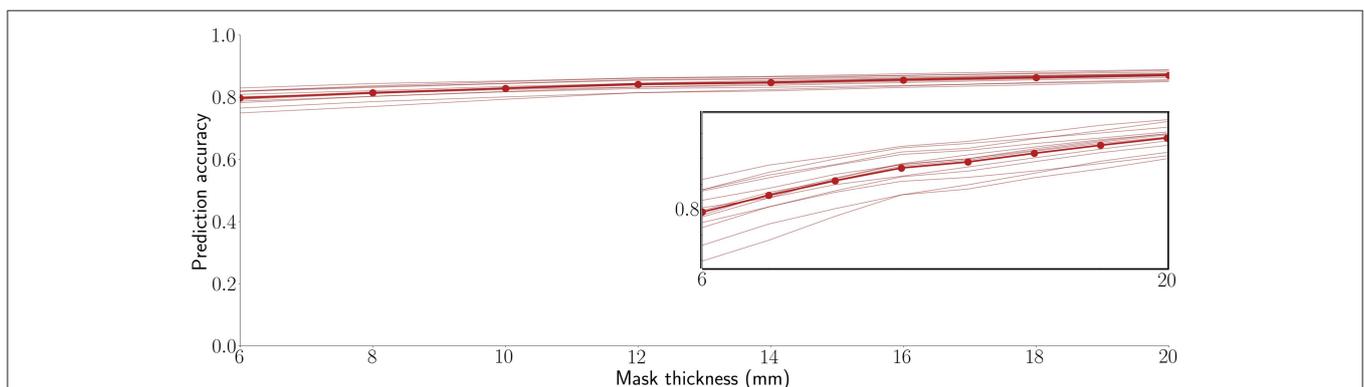
FIGURE 3 | Decoding is highly accurate across a practical range of surface-limited recording depths. Each line is a participant's average across 16 target signals (10 RSNs and 6 anatomical regions). The bold line connected by points shows the average prediction accuracy across all 13 participants. The inset replicates the plot at a magnified vertical scale to emphasize that prediction accuracy increases monotonically with mask thickness and to highlight the observation that, even within the tight range of observed correlation values, participants are highly consistent in their performance ranking across the tested mask thicknesses.
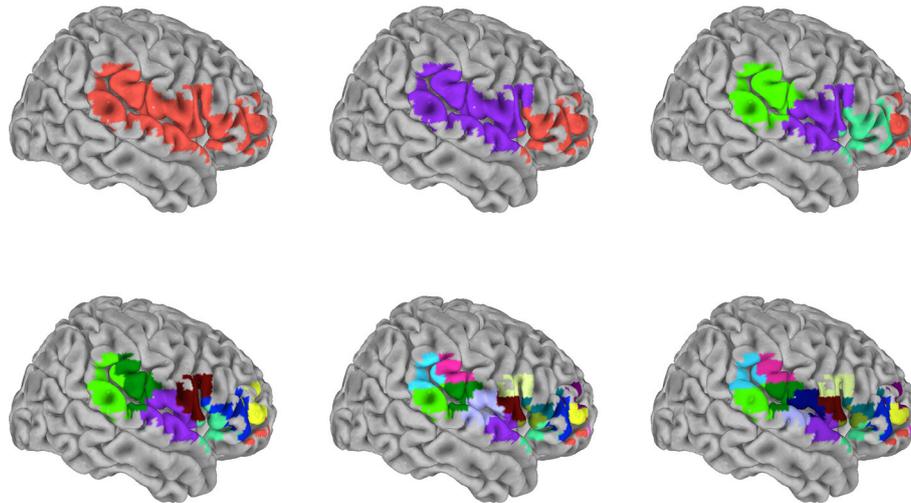
FIGURE 4 | Graphical depiction of the iterative surface volumetric subdivision. Group maps were generated for each RSN and anatomical target signal using the 99-participant SVR surface models. Each significant cluster was then volumetrically subdivided to generate sparse feature sets for each target. This example shows the iterative steps used to subdivide cluster 1 in the insula target group map (see also **Supplementary Table S14**). The cluster spans frontal, postcentral, temporal, and parietal regions. This particular cluster was subdivided in 5 iterations, reducing 489 voxels to 17 subdivisions. Colors distinguish individual subdivisions. Subsequently, the signal average within each subdivision was used as a training/testing feature in the independent 13-participant data.

20). To test these subdivided features, training and test data from the 13-participant set were extracted by averaging within each member voxel of a given subdivision. This was then also examined in the replication data set (openneuro.org; Power et al., 2017).

Decoding Performance, Surface Weight Vectors and Example Target Decoding

Our primary question was, "Can we use the surface of the brain to track distributed networks and anatomical regions?" The prediction accuracies and wavestrap significance thresholds are shown in **Figure 5**, confirming that it is indeed possible to track interior brain regions and distributed RSN fluctuations using only surface measurements (see also **Supplementary Figures S6, S7** for wavestrap details). For consistancy with our mapping results, **Figure 5** uses MNI-normalized data. As **Supplementary Figure S8** shows, however, the prediction accuracies are similar when calculated in in each subject's native (and unsmoothed) space. Moreover, these results were remarkably reproducible in the 83 participants (**Supplementary Figure S9**). Our secondary question was, "Can we map the surface connectivity that enables decoding?" Group maps of the 99 participants' surface weight vectors are shown for the RSN targets in **Figure 6** and for the anatomical targets in **Figure 7**. These maps represent functional connectivity relationships between the surface and each predicted target signal. Specifically, these maps show consistencies in the SVR-derived multivariate combinations of surface time series that best reconstructed the interior brain signal. The maps in **Figures 6, 7** are fully described in **Supplementary Tables S1–S16** and available on neurovault.org (https://neurovault.org/collections/QAGSDTLT/).

In addition, these figures show the best- and worst-predicted run for the 13 participants, corresponding to the full resolution surface (maroon bars) in **Figure 5**. Finally, we examined the residuals of all predictions. Individual examples are include in **Figures 6, 7**. On average, the residuals appear to be normally distributed and absent of strong correlations between the residuals vs. predicted values. See **Supplementary Figures S10, S11**.

4. DISCUSSION

The surface of the brain is comprised of the neocortex, an evolutionarily young, folded sheet of neurons that are the substrate for highly refined sensory, motor, and cognitive function. Its connectivity with itself and the rest of the brain is critical for higher order function and may be one of the keys to understanding neurodegeneration and psychiatric disorders. Importantly, since the cortex forms the exterior of the brain, all but its medial walls are accessible to a broad range of physiological measures. Consequently, however, the surface also physically obscures the deeper structures it surrounds to depth-limited measures. We hypothesized, though, that machine learning models of functional connectivity might enable the surface of the brain to be used as a steerable array, tuned to report activity of deeper regions and distributed networks. If this were true, then potentially even surface-limited measurements could be calibrated to report ongoing internal activity and to examine the multivariate information exchange between the cortex and the rest of the brain.

With this in mind, this study set out to evaluate whether supervised regression models of functional connectivity could harness multivariate signals from the surface of the brain to track
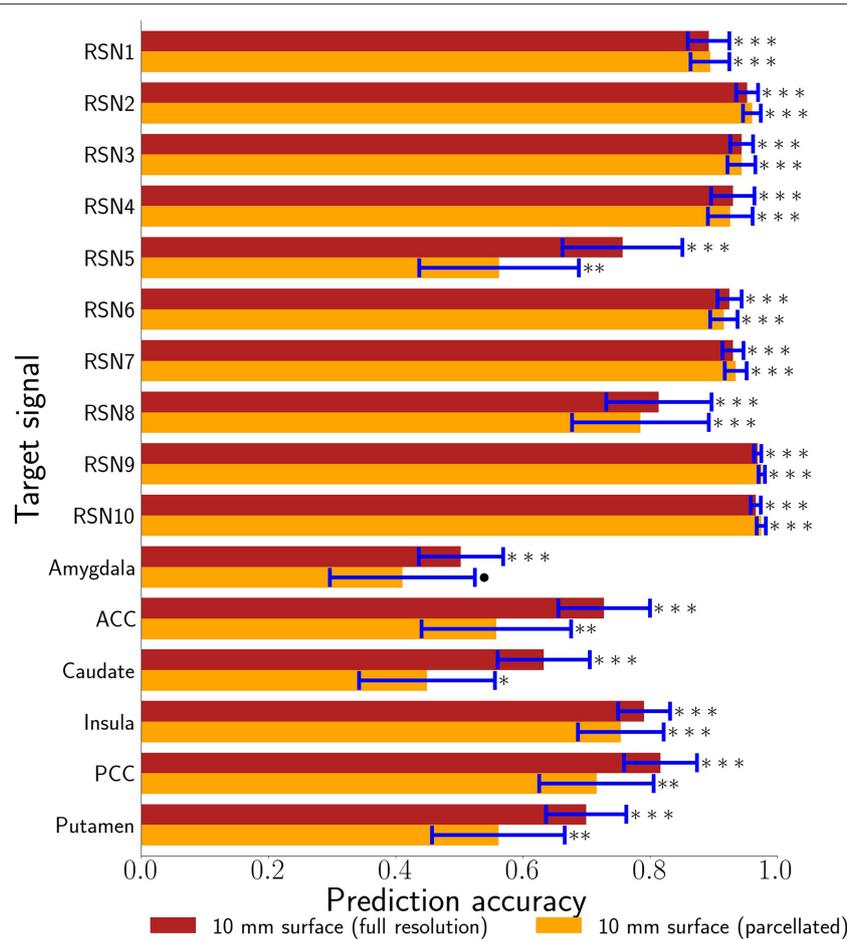
FIGURE 5 | 10 mm depth-limited prediction accuracies for all target signals demonstrates the general feasibility of tracking activity using the brain's surface. The full voxel resolution results are shown in maroon and the volumetrically subdivided results are shown in orange. Error bars are plus-or-minus one standard deviation. Non-parametric estimates of significance were generated from wavestrapped distributions and were corrected for multiple comparisons with respect to these 32 hypothesis tests. *** indicates corrected $p < 0.001$, ** indicates corrected $p < 0.01$, * indicates FDR-corrected $p < 0.05$, ● indicates corrected $p < 0.06$. ACC denotes anterior cingulate cortex. PCC denotes posterior cingulate cortex.

underlying regions and distributed networks. Taken together, the results demonstrate that it is, indeed, possible to decode the brain's surface in such a manner. Thus, while physically surrounding the rest of the brain, the cortex nonetheless remains a window to the deeper activity it encases. These results reveal that a substantial amount of information is shared between the cortex and the rest of the brain through functional connectivity. The fidelity of the decoding depends upon several factors such as the depth of measurement, the specific target network or region, the trade-off between utilizing the entire surface vs. sparse, volumetrically-subdivided features, and the subject-specific variability observed in **Figure 3**. Note also that **Figures 6, 7** indicate subject-specific characteristics for the best and worst predictions. Almost half of the worst performances (7 of 16) come from 2 of the 26 runs: participant 2's first run (ACC, putamen, RSN3, RSN10) and participant 9's s run (insula, RSN5/cerebellum, RSN9). Participant 9 is interesting, however, because while her/his second run represents three of the worst

performances, her/his first run represents three of the best (caudate, RSN3, RSN8). Finally, although not shown here, we and others have previously noted that the amount of training data (in this case, the length of the fMRI runs) also plays an important role in determining the prediction accuracy of classifiers and the reliability of estimates (Kjems et al., 2002; LaConte et al., 2007; Pereira et al., 2009).

The RSN accuracies were generally higher than those for the bilateral anatomical regions included in this study. The results corroborate that these networks can, indeed, be treated as a unit and that the functional coherence that originally defined them (Smith et al., 2009) is replicated here in a new cohort and in a new analysis context. One important factor that also likely contributed to the strong performance of the RSNs is that their spatial extent partially comprises the surface, itself. Thus, there is a redundancy between the RSN voxels averaged to define the target signal and the surface features, vastly simplifying the training complexity. The exception is the cerebellar network
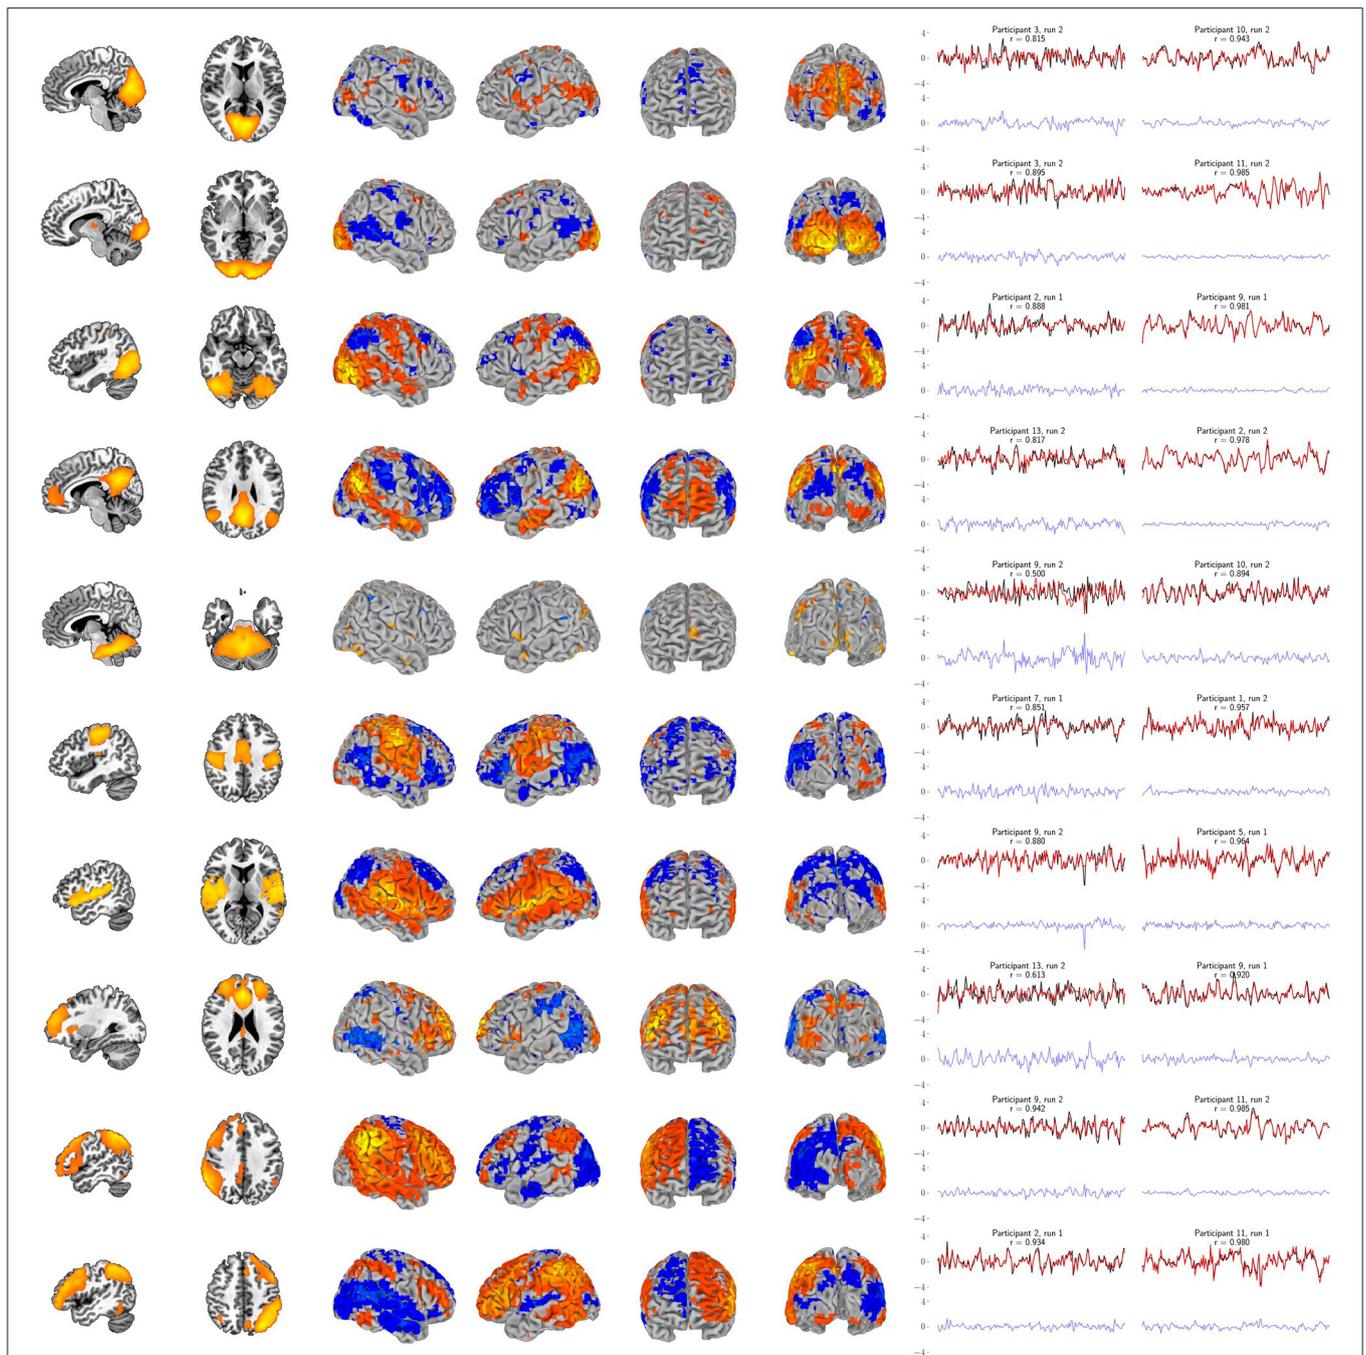
FIGURE 6 | Support vector regression group maps for the ten resting-state networks from Smith et al. (2009). Left panel: The resting-state network templates (with RSN1 to RSN10 ordered from top to bottom, respectively). Center panel: Group maps from the 99-participant SVR models using the full-resolution (10 mm depth) mask for each participant (FDR-corrected $p < 0.05$). Group maps for all network targets are fully described in **Supplementary Tables S1–S10**. Right panel: The least and most accurate predictions for each target, respectively from the 13-participant data set. The actual target time series is black and the surface-limited prediction is red. Residual time series (target - predicted) are shown on the same scale below in blue. Annotations indicate the participant number, resting-state run number and correlation (r) between the target and predicted time series. For example, "Participant 1, run 2" indicates that a surface SVR model trained on that participant's run 1 data was used to decode her/his run 2 target activity (red time series). See **Figure 7** for corresponding results for the six bilateral anatomical targets.

(RSN5), which is also notable in that it demonstrated the largest performance drop when going from the full surface model to the volumetrically subdivided one. This suggests that the 99-subject group model does not generalize as well for RSN5 and that more extensive surface coverage aids in predicting cerebellum when it is defined as a cohesive network. Future

FIGURE 7 | Support vector regression for the six bilateral anatomical targets. Left panel: Bilateral anatomical target regions. Ordered from top to bottom are amygdala, anterior cingulate, caudate, insula, posterior cingulate, and putamen. Center panel: Group maps from the 99-participant SVR models using the full-resolution (10 mm depth) mask for each participant (FDR-corrected $p < 0.05$). Group maps for all anatomical targets are fully described in **Supplementary Tables S11–S16**. Right panel: The least and most accurate predictions for each target, respectively from the 13-participant data set. The actual target time series is black and the surface-limited prediction is red. Residual time series (target-predicted) are shown on the same scale below in blue. Annotations indicate the participant number, resting-state run number and correlation (r) between the target and predicted time series. For example, "Participant 1, run 2" indicates that a surface SVR model trained on that participant's run 1 data was used to decode her/his run 2 target activity (red time series). See **Figure 6** for corresponding RSN results.

studies are needed to further characterize functional connectivity of the cortex with anatomical and functional sub-regions of the cerebellum. Similar considerations are warranted for the six bi-lateral anatomical regions. It is possible that their accuracy might be improved by modeling smaller, unilateral subregions. From a signal-to-noise point of view, it might be ideal to define target signals by averaging across regions or networks that have optimized the statistical trade-offs between maximizing spatial extent and maximizing signal homogeneity (Shirer et al., 2012). On the other hand, mixtures of signal sources may or may not degrade accuracy, depending upon how well this mixture can be modeled from the surface features. Thus from a more nuanced view, both the bi-lateral anatomical regions and the RSNs used in this study can be considered as networks. Further, it is important to recall that fMRI voxels are several orders of magnitude larger than their underlying cellular processes and are thus fundamentally mixtures of signal sources. Related to these issues, further improvements could possibly come from adaptations of hyperalignment (Haxby et al., 2011; Feilong et al., 2021), which could be used to account for individual variations in coarse-grained patterns and to maximize the similarity within

ROIs and networks as well as with the brain's surface. Overall, we expect that tracking both smaller subdivisions as well as extensive distributed networks will be important in future research. Additional studies will be needed to comprehensively test the factors that determine how well any given target signal can be tracked from the surface. To summarize, all regions and networks are a mixture of several sources and prediction accuracy is a metric to evaluate how functionally connected to the cortex the chosen target signal is. At present, the broad success of all sixteen target signals and thus the conclusion that these are all simultaneously represented at the level of surface cortex is fascinating.

fMRI proved to be ideal for demonstrating the feasibility of surface-based tracking. Because of its high resolution, it enabled examination of the extremes of surface density measurements to evaluate the theoretical impact of high and low numbers of training features. And in the majority of cases, high accuracy was maintained with a much smaller feature space. In addition, validation is possible because the fMRI data are whole-brain. The two runs in the 13-participant data set and in the 83-participant replication set allowed us to estimate

our primary measure, surface-based decoding accuracy, in a cross-validation framework (alternating the role of training and testing runs and averaging the two prediction accuracy estimates). For today's technology, the high accuracy reported here likely serves as an optimistic upper bound on what could be accomplished in an actual multimodal setting. As pointed out by the reviewers of this work, spatial correlations, head motion, global signal, and physiology are all sources of noise that may need to be further investigated. Moreover, the relative impact of each of these will almost certainly vary from modality to modality, and thus may need different strategies. One interesting possibility, however, is that variations on our approach may lead to new ways to incorporate such nuisance signals directly into the multi-modal modeling approach. As shown in **Supplementary Figures S10–S11**, the residuals from our surface models appear to be largely white noise—notably unlike any of the confounding potential noise sources that we have listed. Thus, in our case, the preprocessing that we applied combined with the supervised learning process seemed to effectively regress out these confounding signals. Another point that should be elaborated upon is that the goal of our study deviates from what some would consider usual notions of data independence. For example, many of the RSNs extend to the brain's surface. Rather than excluding surface data features, we actually want them to be incorporated into the training stage. As shown in **Figure 6** these surface portions of the networks are highly likely to be important features in the model. But what **Figure 6** also shows is that the models do not exclusively rely on such regions. Indeed if they did, that would be a sign that just that region and not its *network* was being disproportionately modeled. As **Figure 6** shows, this is not what is observed. Even RSNs 1–3, which are close to the occipital pole, converge to highly distributed surface patterns. Finally, along these lines, partial volume effects for fMRI data could partially serve as an aid or confound for tracking regions with surface measures. Analogous to partial volume concerns, each surface modality has its own individual depth sensitivity profile. Thus, while this study demonstrates the feasibility of using surface models to track deeper regions and networks, future work will be needed to address specific concerns for adopting this for actual surface recordings, and validation of such future work will likely require simultaneous multimodal measurements. Finally, while our stated goal was to demonstrate a simple approach to surface modeling and specifically to preserve the folded surface, it is possible that future improvements could come from more sophisticated techniques. As the reviewers have pointed out, with volume-based registration brain voxels around the cortical surface can suffer from severe misalignment issues. While a careful comparison would constitute a study of its own, it is possible that surface-based registration may be a necessary future step for future multi-modal implementations. At this point, however, our intuition is that challenges with surface probe/sensor placement will constitute a larger source of future variance.

In a complementary way, the 99-participant data set provided ample training data to examine commonalities of the SVR models across individuals. As the maps show (center panels in

Figures **6**, **7**), this analysis approach has a secondary benefit as a tool for studying functional connectivity with the cortex. And, since the volumetrically-subdivided masks that these group maps generated were independent from the 13- and 83-participant sets, we also had the opportunity to go beyond mapping and validate the predictive potential of these functional connectivity findings. On the whole, the group maps representing the surface models of the 6 anatomical regions and the 10 RSNs are highly informative. Since the cortex's functional connectivity is thought to be one of the enabling properties of higher order function, this methodological approach could provide a new avenue for precisely exploring specific, restricted connectivity questions.

For example, the amygdala's functional roles and its connectivity with the brain's surface are continually being refined. The results in **Figure 7** (see also **Supplementary Table S11**) overlap with frontal and temporal regions reported by Bickart et al. (2012, 2014), who explore this region's role as a critical hub of the social brain. This is just one recent case illustrating how our understanding of the function of the amygdala and its connectivity with the brain's surface continues to evolve. The first comprehensive description of the cortical projections of the amygdala in primates was reported by Amaral and Price (1984), who also described the early history uncovering its cortical projections. As they note, amygdalo-cortical connections were neither widely known nor appreciated until the early 1960s. As the complexity of the amygdala's connectivity was gradually discovered, so too was the recognition that its functional role was also likely to be highly complex. Thus, while our driving question was "Can limited surface models be used to track networks and deeper brain regions?," our approach also doubles as a mapping tool to explore functional connectivity with the brain's surface.

Our analysis treats the cortex as a tunable receiver array, but this conceptual framework should not be taken literally, and it is important to note that the direction of information exchange between the surface and other brain regions cannot be resolved with our approach. It is likely that these models reflect both the cortex's signals to and from other regions. The ability to examine functional connectivity and to use this property to track time series derives directly from previous work on multivariate functional integration (Friston et al., 1993; Friston, 1994; Chu et al., 2011; Craddock et al., 2013). This approach has been applied to examine functional hierarchies using both connectivity maps and prediction accuracy (Craddock et al., 2013) through the "non-parametric prediction accuracy, influence, and reproducibility resampling" (NPAIRS) framework (Strother et al., 2002; LaConte et al., 2003). It is important to note that our previous results (Craddock et al., 2013) are not directly comparable to the those reported here. In Craddock et al. (2013) we used parcellated features throughout the entire brain, including subcortical areas, while this study used features restricted to the most superficial surface of the brain. So, for example, our full-resolution amygdala results have an average prediction accuracy of $r = 0.5$ and our sparse surface features give $r = 0.41$, while the subjects in Craddock et al. (2013) average about $r = 0.8$. This reflects expected decreases in accuracy because much of the amygdala's

functional connectivity is unavailable to the surface-restricted models explored here.

Additional studies of functional integration include Hay et al. (2017) who demonstrated that minimal sets of features could be derived with recursive feature elimination and used to quantify regional degree of functional integration. Further, multivariate functional integration model weights have been demonstrated as a means of connectotyping (or functionally fingerprinting) individuals since they produce a personalized model-based connectivity matrix (Miranda-Dominguez et al., 2014). In this current study we have demonstrated the utility of restricting the functional integration investigation through the use of surface training features. Several additional variations of this approach are possible to study the functional integration between combinations of functional or anatomical nodes and networks which would complement graph theoretic (Bullmore and Sporns, 2009) and dynamic causal modeling (Friston et al., 2003) studies. Though not explored here, the utility of functional integration may be enhanced by comparing and contrasting different loss functions provided by alternative regression techniques such as LASSO (Tibshirani, 1996) and elastic nets (Zou and Hastie, 2005).

Restricting the training features to the surface of the brain enabled us to computationally assess the feasibility of multimodal enhancements to depth-limited measurements. Three major factors interplay when considering the possibility of using surface recordings to reconstruct deeper activity. The first (which was the primary focus of this work) concerns algorithmic and data properties, such as the density of measurements and their distributed coverage. A common goal in machine learning is to capture as much information as possible with the fewest possible data features. For practical reasons, it is desirable to assess the impact of sparse recordings as well as the impact on depth of sensitivity. fMRI is well-suited to study these issues since it provides whole brain coverage at a relatively high spatial resolution, surpassing what is currently achievable in terms of both density and depth of surface recordings. The second factor is the structure of the cortex, itself. Although its flattened dimension is approximately 2.5 mm thick on average, ranging from 1 to 4.5 mm (Fischl and Dale, 2000), its folded sulcal depths in humans have been reported to range from 1 to 3 cm (Ribas, 2010). Thus, the relative distance and orientation of gray matter varies dramatically as a function of position along the surface of the brain. Looking forward to future studies, the third factor is that each surface measurement has unique physical limitations. For example, MEG records magnetic fields outside the head that are generated by neural currents. Its sensitivity depends on both the depth and orientation of the generating current sources (Goldenholz et al., 2009). In contrast, functional near-infrared spectroscopy (fNIRS) uses two near-infrared wavelengths to assess changes in blood oxygenation arising from neural metabolic demands. Estimates of fNIRS penetration depth varies in the literature, ranging from 12 to 21 mm into the brain (Okamoto et al., 2004; Schroeter et al., 2006; Lu et al., 2010).

Moving beyond the feasibility of surface-limited tracking demonstrated here, implementation could take several forms. The basic notion is that a modality such as fMRI (that can measure both target and surface training signals) be used to calibrate a surface-limited modality such as fNIRS, MEG, or EEG. Such a process would enable those modalities to track activity that far surpasses their inherent physical capabilities. Early empirical work suggests that this is possible in EEG, in which amygdala activity can be tracked using an EEG fingerprint (Meir-Hasson et al., 2014) and in fNIRS, where similar analyses confirmed the ability to track deep brain regions in task-based fMRI (Liu et al., 2015). Thus, there is early positive evidence suggesting that it is possible to dramatically enhance surface limited methods well beyond their intrinsic physical capabilities. In the meantime, we have additionally shown that the brain has a rich surface connectivity that can be mapped and has simultaneous predictive potential for decoding a range of distributed networks as well as interior, localized regions.

DATA AVAILABILITY STATEMENT

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found below: https://neurovault.org/collections/QAGSDTLT/.

ETHICS STATEMENT

The studies involving human participants were reviewed and approved by Virginia Tech IRB. The patients/participants provided their written informed consent to participate in this study.

AUTHOR CONTRIBUTIONS

MT: investigation, methodology, software, and writing. JL: investigation, methodology, software, data curation, and writing. SL: investigation, methodology, supervision, and writing. All authors contributed to the article and approved the submitted version.

FUNDING

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fnimg.2022.815778/full#supplementary-material

REFERENCES

Amaral, D. G., and Price, J. L. (1984). Amygdalo-cortical projections in the monkey (*Macaca fascicularis*). *J. Comp. Neurol.* 230, 465–496. doi: 10.1002/cne.902300402

Bickart, K. C., Dickerson, B. C., and Barrett, L. F. (2014). The amygdala as a hub in brain networks that support social life. *Neuropsychologia* 63, 235–248. doi: 10.1016/j.neuropsychologia.2014.08.013

Bickart, K. C., Hollenbeck, M. C., Barrett, L. F., and Dickerson, B. C. (2012). Intrinsic amygdala-cortical functional connectivity predicts social network size in humans. *J. Neurosci.* 32, 14729–14741. doi: 10.1523/JNEUROSCI.1599-12.2012

Breakspear, M., Brammer, M. J., Bullmore, E. T., Das, P., and Williams, L. M. (2004). Spatiotemporal wavelet resampling for functional neuroimaging data. *Hum. Brain Mapp.* 23, 1–25. doi: 10.1002/hbm.20045

Bullmore, E., Fadili, J., Maxim, V., Şendur, L., Whitcher, B., Suckling, J., et al. (2004). Wavelets and functional magnetic resonance imaging of the human brain. *Neuroimage* 23, S234–S249. doi: 10.1016/j.neuroimage.2004.07.012

Bullmore, E., and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.* 10, 186–198. doi: 10.1038/nrn2575

Chu, C., Handwerker, D. A., Bandettini, P. A., and Ashburner, J. (2011). "Measuring the consistency of global functional connectivity using kernel regression methods," in *2011 International Workshop on Pattern Recognition in NeuroImaging* (Seoul), 41–44.

Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput. Biomed. Res.* 29, 162–173. doi: 10.1006/cbmr.1996.0014

Craddock, C., Milham, M., and LaConte, S. (2013). Predicting intrinsic brain activity. *Neuroimage* 82, 127–136. doi: 10.1016/j.neuroimage.2013.05.072

Dale, A. M., Liu, A. K., Fischl, B. R., Buckner, R. L., Belliveau, J. W., Lewine, J. D., et al. (2000). Dynamic statistical parametric mapping: combining fMRI and MEG for high-resolution imaging of cortical activity. *Neuron* 26, 55–67. doi: 10.1016/S0896-6273(00)81138-1

Eickhoff, S. B., Stephan, K. E., Mohlberg, H., Grefkes, C., Fink, G. R., Amunts, K., et al. (2005). A new SPM toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data. *Neuroimage* 25, 1325–1335. doi: 10.1016/j.neuroimage.2004.12.034

Feilong, M., Guntupalli, J. S., and Haxby, J. V. (2021). The neural basis of intelligence in fine-grained cortical topographies. *Elife* 10, e64058. doi: 10.7554/eLife.64058

Feinberg, D. A., Moeller, S., Smith, S. M., Auerbach, E., Ramanna, S., Glasser, M. F., et al. (2010). Multiplexed echo planar imaging for sub-second whole brain FMRI and fast diffusion imaging. *PLoS ONE* 5, e15710. doi: 10.1371/journal.pone.0015710

Fiederer, L., Vorwerk, J., Lucka, F., Dannhauer, M., Yang, S., Dümpelmann, M., et al. (2016). The role of blood vessels in high-resolution volume conductor head modeling of EEG. *NeuroImage*, 128(Suppl. C), 193–208. doi: 10.1016/j.neuroimage.2015.12.041

Fischl, B. (2012). Freesurfer. *NeuroImage* 62, 774–781. doi: 10.1016/j.neuroimage.2012.01.021

Fischl, B., and Dale, A. M. (2000). Measuring the thickness of the human cerebral cortex from magnetic resonance images. *Proc. Natl. Acad. Sci. U.S.A.* 97, 11050–11055. doi: 10.1073/pnas.200033797

Friston, K., Harrison, L., and Penny, W. (2003). Dynamic causal modelling. *Neuroimage* 19, 1273–1302. doi: 10.1016/S1053-8119(03)00202-7

Friston, K. J. (1994). Functional and effective connectivity in neuroimaging: a synthesis. *Hum. Brain Mapp.* 2, 56–78. doi: 10.1002/hbm.460020107

Friston, K. J., Frith, C. D., and Frackowiak, R. S. J. (1993). Time-dependent changes in effective connectivity measured with PET. *Hum. Brain Mapp.* 1, 69–79. doi: 10.1002/hbm.460010108

Friston, K. J., Williams, S., Howard, R., Frackowiak, R. S. J., and Turner, R. (1996). Movement-related effects in fMRI time-series. *Mag. Reson. Med.* 35, 346–355. doi: 10.1002/mrm.1910350312

Goldenholz, D. M., Ahlfors, S. P., Hämäläinen, M. S., Sharon, D., Ishitobi, M., Vaina, L. M., et al. (2009). Mapping the signal-to-noise-ratios of cortical sources in magnetoencephalography and electroencephalography. *Hum. Brain Mapp.* 30, 1077–1086. doi: 10.1002/hbm.20571

Griswold, M. A., Jakob, P. M., Heidemann, R. M., Nittka, M., Jellus, V., Wang, J., et al. (2002). Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Mag. Reson. Med.* 47, 1202–1210. doi: 10.1002/mrm.10171

Haxby, J. V., Guntupalli, J. S., Connolly, A. C., Halchenko, Y. O., Conroy, B. R., Gobbini, M. I., et al. (2011). A common, high-dimensional model of the representational space in human ventral temporal cortex. *Neuron* 72, 404–416. doi: 10.1016/j.neuron.2011.08.026

Hay, E., Ritter, P., Lobaugh, N. J., and McIntosh, A. R. (2017). Multiregional integration in the brain during resting-state fMRI activity. *PLoS Comput. Biol.* 13, e1005410. doi: 10.1371/journal.pcbi.1005410

Jenkinson, M., Bannister, P., Brady, M., and Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* 17, 825–841. doi: 10.1006/nimg.2002.1132

Joachims, T. (1999). "Making large-scale SVM learning practical," in *Advances in Kernel Methods-Support Vector Learning, Chapter 11*, eds B. Schölkopf, C. Burges, and A. Smola (Cambridge, MA: MIT Press), 169–184.

Kjems, U., Hansen, L., Anderson, J., Frutiger, S., Muley, S., Sidtis, J., et al. (2002). The quantitative evaluation of functional neuroimaging experiments: mutual information learning curves. *Neuroimage* 15, 772–786. doi: 10.1006/nimg.2001.1033

LaConte, S., Anderson, J., Muley, S., Ashe, J., Frutiger, S., Rehm, K., et al. (2003). The evaluation of preprocessing choices in single-subject BOLD fMRI using NPAIRS performance metrics. *Neuroimage* 18, 10–27. doi: 10.1006/nimg.2002.1300

LaConte, S., Strother, S., Cherkassky, V., Anderson, J., and Hu, X. (2005). Support vector machines for temporal classification of block design fMRI data. *Neuroimage* 26, 317–329. doi: 10.1016/j.neuroimage.2005.01.048

LaConte, S. M., Peltier, S. J., and Hu, X. P. (2007). Real-time fMRI using brain-state classification. *Hum. Brain Mapp.* 28, 1033–1044. doi: 10.1002/hbm.20326

Larkman, D. J., Hajnal, J. V., Herlihy, A. H., Coutts, G. A., Young, I. R., and Ehnholm, G. (2001). Use of multicoil arrays for separation of signal from multiple slices simultaneously excited. *J. Mag. Reson. Imaging* 13, 313–317. doi: 10.1002/1522-2586(200102)13:2andlt;313::AID-JMRI1045andgt;3.0.CO;2-W

Lauterbur, P. C. (1973). Image formation by induced local interactions: Examples employing nuclear magnetic resonance. *Nature* 242, 190–191. doi: 10.1038/242190a0

Liu, N., Cui, X., Bryant, D. M., Glover, G. H., and Reiss, A. L. (2015). Inferring deep-brain activity from cortical activity using functional near-infrared spectroscopy. *Biomed. Opt. Express* 6, 1074–1089. doi: 10.1364/BOE.6.001074

Lu, C.-M., Zhang, Y.-J., Biswal, B. B., Zang, Y.-F., Peng, D.-L., and Zhu, C.-Z. (2010). Use of fNIRS to assess resting state functional connectivity. *J. Neurosci. Methods* 186, 242–249. doi: 10.1016/j.jneumeth.2009.11.010

Lund, T. E., Madsen, K. H., Sidaros, K., Luo, W.-L., and Nichols, T. E. (2006). Non-white noise in fMRI: does modelling have an impact? *Neuroimage* 29, 54–66. doi: 10.1016/j.neuroimage.2005.07.005

Mansfield, P. (1977). Multi-planar image formation using nmr spin echoes. *J. Phys. C* 10, L55. doi: 10.1088/0022-3719/10/3/004

Meir-Hasson, Y., Kinreich, S., Podlipsky, I., Hendler, T., and Intrator, N. (2014). An EEG finger-print of fMRI deep regional activation. *Neuroimage* 102, 128–141. doi: 10.1016/j.neuroimage.2013.11.004

Menon, V., Ford, J. M., Lim, K. O., Glover, G. H., and Pfefferbaum, A. (1997). Combined event-related fMRI and EEG evidence for temporal-parietal cortex activation during target detection. *Neuroreport* 8, 3029–3037. doi: 10.1097/00001756-199709290-00007

Meyer-Lindenberg, A. (2010). From maps to mechanisms through neuroimaging of schizophrenia. *Nature* 468, 194–202. doi: 10.1038/nature09569

Miranda-Dominguez, O., Mills, B. D., Carpenter, S. D., Grant, K. A., Kroenke, C. D., Nigg, J. T., et al. (2014). Connectotyping: model based fingerprinting of the functional connectome. *PLoS ONE* 9, e0111048. doi: 10.1371/journal.pone.0111048

Moeller, S., Yacoub, E., Olman, C. A., Auerbach, E., Strupp, J., Harel, N., et al. (2010). Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal

whole-brain fMRI. *Mag. Reson. Med.* 63, 1144–1153. doi: 10.1002/mrm.22361

Mugler, J. P., and Brookeman, J. R. (1990). Three-dimensional magnetization-prepared rapid gradient-echo imaging (3D MP RAGE). *Mag. Reson. Med.* 15, 152–157. doi: 10.1002/mrm.1910150117

Nunes, R., Hajnal, J., Golay, X., and Larkman, D. (2006). "Simultaneous slice excitation and reconstruction for single shot EPI," in *Proceedings of the 14th Annual Meeting of ISMRM* (Seattle, WA), 293.

Okamoto, M., Dan, H., Sakamoto, K., Takeo, K., Shimizu, K., Kohno, S., et al. (2004). Three-dimensional probabilistic anatomical cranio-cerebral correlation via the international 10-20 system oriented for transcranial functional brain mapping. *Neuroimage* 21, 99–111. doi: 10.1016/j.neuroimage.2003.08.026

Pascual-Marqui, R. D. (1999). Review of methods for solving the EEG inverse problem. *Int. J. Bioelectromagn.* 1, 75–86.

Pereira, F., Mitchell, T., and Botvinick, M. (2009). Machine learning classifiers and fMRI: a tutorial overview. *Neuroimage* 45(Suppl. 1), S199–S209. doi: 10.1016/j.neuroimage.2008.11.007

Power, J. D., Plitt, M., Kundu, P., Bandettini, P. A., and Martin, A. (2017). Temporal interpolation alters motion in fMRI scans: magnitudes and consequences for artifact detection. *PLoS ONE* 12, e0182939. doi: 10.1371/journal.pone.0182939

Pruessmann, K. P., Weiger, M., Scheidegger, M. B., and Boesiger, P. (1999). SENSE: sensitivity encoding for fast MRI. *Mag. Reson. Med.* 42, 952–962. doi: 10.1002/(SICI)1522-2594(199911)42:5<952::AID-MRM16>3.0.CO;2-S

Ribas, G. C. (2010). The cerebral sulci and gyri. *Neurosurg. Focus* 28, E2. doi: 10.3171/2009.11.FOCUS09245

Sato, J. R., Rondinoni, C., Sturzbecher, M., de Araujo, D. B., and Amaro Jr, E. (2010). From eeg to bold: brain mapping and estimating transfer functions in simultaneous eeg-fmri acquisitions. *Neuroimage* 50, 1416–1426. doi: 10.1016/j.neuroimage.2010.01.075

Schroeter, M. L., Kupka, T., Mildner, T., Uludağ, K., and von Cramon, D. Y. (2006). Investigating the post-stimulus undershoot of the BOLD signal-a simultaneous fMRI and fNIRS study. *Neuroimage* 30, 349–358. doi: 10.1016/j.neuroimage.2005.09.048

Shirer, W. R., Ryali, S., Rykhlevskaia, E., Menon, V., and Greicius, M. D. (2012). Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cereb. Cortex* 22, 158–165. doi: 10.1093/cercor/bhr099

Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., et al. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci. U.S.A.* 106, 13040–13045. doi: 10.1073/pnas.0905267106

Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E., Johansen-Berg, H., et al. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* 23(Suppl. 1), S208–S219. doi: 10.1016/j.neuroimage.2004.07.051

Sodickson, D. K., and Manning, W. J. (1997). Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Mag. Reson. Med.* 38, 591–603. doi: 10.1002/mrm.1910380414

Strother, S. C., Anderson, J., Hansen, L. K., Kjems, U., Kustra, R., Sidtis, J., et al. (2002). The quantitative evaluation of functional neuroimaging experiments: The NPAIRS data analysis framework. *Neuroimage* 15, 747–771. doi: 10.1006/nimg.2001.1034

Tange, O. (2011). GNU parallel-the command-line power tool. *USENIX Mag.* 36, 42–47.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. B* 58, 267–288. doi: 10.1111/j.2517-6161.1996.tb02080.x

Wasilewski, F. (2006). *PyWavelets: Discrete Wavelet Transform in Python*. Available online at: http://www.pybytes.com/pywavelets/.

Zhang, Y., Brady, M., and Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans. Med. Imaging* 20, 45–57. doi: 10.1109/42.906424

Zou, H., and Hastie, T. (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. B* 67, 301–320. doi: 10.1111/j.1467-9868.2005.00503.x