

Mapping Brain-Behavior Space Relationships Along the Psychosis Spectrum

Jie Lisa Ji^{1,2}, Markus Helmer¹, Clara Fonteneau¹, Joshua B. Burt^{1,3}, Zailyn Tamayo¹, Jure Demšar⁴, Brendan Adkinson^{1,2}, Aleksandar Savić⁵, Katrin H. Preller⁶, Flora Moujaes⁶, Franz X. Vollenweider⁶, William J. Martin⁷, Grega Repovš⁴, John D. Murray^{1,2,3}, and Alan Anticevic^{1,2,8}✉

¹Department of Psychiatry, Yale University School of Medicine, New Haven, CT, USA

²Interdepartmental Neuroscience Program, Yale University School of Medicine, New Haven, CT, USA

³Department of Physics, Yale University, New Haven, CT, USA

⁴Department of Psychology, University of Ljubljana, Ljubljana, Slovenia

⁵Department of Psychiatry, University of Zagreb, Zagreb, Croatia

⁶University Hospital for Psychiatry Zurich, Zurich, Switzerland

⁷The Janssen Pharmaceutical Companies of Johnson & Johnson

⁸Department of Psychology, Yale University School of Medicine, New Haven, CT, USA

Difficulties in advancing effective patient-specific therapies for psychiatric disorders highlights a need to develop a neurobiologically-grounded, quantitatively stable mapping between neural and symptom variation. This gap is particularly acute for psychosis-spectrum disorders (PSD). Here, in a sample of 436 cross-diagnostic PSD patients, we derived and replicated a data-driven dimensionality-reduced symptom space across hallmark psychopathology symptoms and cognitive deficits, which was predictive at the single patient level. In turn, these data-reduced symptom axes mapped onto distinct and replicable univariate brain maps. Critically, we found that multivariate brain-behavior mapping techniques (e.g. canonical correlation analysis) did not show stable results. Instead, we show that a univariate brain-behavioral space (BBS) mapping can resolve stable individualized prediction. Finally, we show a proof-of-principle framework for relating personalized BBS metrics with molecular targets via serotonin and glutamate receptor manipulations and gene expression maps. Collectively, these results highlight a stable and data-driven BBS mapping across PSD, which offers an actionable quantitative path that can be iteratively optimized for personalized clinical biomarker endpoints.

psychosis spectrum | functional neuroimaging | symptoms mapping | molecular target selection | personalized patient selection | statistical learning | data reduction

Correspondence: alan.anticevic@yale.edu & jielisa.ji@yale.edu

Introduction

Mental health conditions cause profound disability with most treatments yielding limited efficacy across psychiatric symptoms (1–4). A key step towards developing more effective therapies for specific psychiatric symptoms is reliably mapping them onto underlying neural systems. This goal is challenging because neuropsychiatric diagnostics still operate under “legacy” categorical constraints, which were not informed by quantitative neural or symptom data.

Critically, diagnostic systems in psychiatry, such as the Diagnostic and Statistical Manual of Mental Disorders (DSM) (5), were built to aid clinical consensus. However, they were not designed to guide quantitative mapping of symptoms onto neural alterations (6, 7). Consequently, the current diagnostic framework cannot, by definition, optimally map onto

patient-specific brain-behavioral alterations. This challenge is particularly evident along the psychosis spectrum disorders (PSD) where there is notable symptom variation across what DSM considers distinct diagnostic categories (such as schizophrenia (SZP), schizo-affective (SADP), bipolar disorder with psychosis (BPP)). For instance, despite BPP being a distinct DSM diagnosis, BPP patients exhibit similar, but attenuated psychosis symptoms and neural alterations similar to SZP (e.g. thalamic functional connectivity (FC) (8)). It is essential to quantitatively map such shared clinical variation onto common neural alterations to circumvent constraints for biomarker development (6, 7, 9) – a key goal for development of neurobiologically-informed personalized therapies (8, 9).

Recognizing the limits of categorical frameworks the NIMH’s Research Domain Criteria (RDoC) initiative introduced dimensional mapping of “functional domains” on to neural circuits (10). This motivated cross-diagnostic multi-site studies to map PSD symptom and neural variation (11–14). In turn, multivariate neuro-behavioral analysis across PSD and mood spectra observed brain-behavioral relationships across diagnoses, with the goal of informing individualized treatment (15). A key challenge that these studies attempted to address is to move beyond *a priori* clinical scales, which provide composite scores (16) that may not optimally capture neural variation (7, 17). For instance, despite many data-driven dimensionality reduction symptom studies (18–24), a common approach in PSD neural research is still to sum “positive” or “negative” psychosis symptoms into a single score for mapping onto neural features. Importantly, altered symptom-to-neural variation may reflect a more complex weighted symptom combination beyond such composite scores.

As noted, multivariate neuro-behavioral studies attempted to address this, but have failed to replicate due to overfitting resulting from high dimensionality of behavioral and neural features (25). It is possible that a linearly-weighted lower-dimensional symptom solution (capturing key disease-relevant information) produces a replicable and robust univariate brain-behavioral mapping. Indeed, recent work used dimensionality reduction methods successfully to compute a

neural mapping across canonical SZP symptoms (23). However, it is unknown if this approach generalizes across PSD. Moreover, it is unknown if incorporating cognitive assessment, a hallmark and untreated PSD deficit (17), explains neural feature variation that is distinct from canonical PSD symptoms. Finally, prior work has not tested if a low-dimensional symptom-to-neural mapping can be predictive at the single patient level – a prerequisite for individualized clinical endpoints.

To inform these gaps, we first tested two key questions: i) Can data-reduction methods reliably map symptom axes across PSD that include both canonical symptoms and cognitive deficits? ii) Do these lower-dimensional symptom axes map onto a replicable brain-behavioral solution across PSD? Specifically, we combined fMRI-derived resting-state measures with psychosis and cognitive symptoms (26, 27) obtained from a public multi-site cohort of 436 PSD patients and 202 healthy individuals collected by the Bipolar-Schizophrenia Network for Intermediate Phenotypes (BSNIP-1) consortium across 6 sites in North America (11). The dataset included 150 patients formally diagnosed with BPP, 119 patients diagnosed SADP, and 167 patients diagnosed with SZP (Table S1). This cohort enabled symptom-to-neural cross-site comparisons across multiple psychiatric diagnostic categories, which we then mapped onto specific neural circuits. We tested if PSD symptoms map onto a low-dimensional solution that is stable for individual patient prediction. Next, we tested if this low-dimensional symptom solution yields novel and replicable neural mapping compared to canonical psychosis composite scores or DSM diagnoses. In turn, we tested if the computed symptom-to-neural mapping is replicable across symptom axes and actionable for individual patient prediction. Finally, we ‘benchmark’ the derived symptom-relevant neural maps by computing their similarity against independently collected pharmacological fMRI maps from healthy adults in response to putative PSD receptor treatment targets (glutamate via ketamine and serotonin via LSD) (28, 29). In turn, we used the Allen Human Brain Atlas (AHBA) to compute gene expression maps (30, 31) for targets implicated in PSD (i.e. interneurons, serotonin and GABA receptor genes). In turn, we tested if the gene targets map onto the derived symptom-relevant neural targets. Collectively, this study used data-driven dimensionality reduction methods to map orthogonal and stable symptom dimensions across 436 PSD patients. In turn, the goal was to map novel and replicable symptom-to-neural relationships across the PSD derived from a low-dimensional symptom solution. Finally, these effects were benchmarked against molecular imaging targets that may be actionable for individualized clinical endpoints. An overview is shown in Fig. S1.

Results

Dimensionality-Reduced PSD Symptom Variation is Stable and Replicable. First, to evaluate PSD symptom variation we examined two instruments: the Brief Assessment of Cognition in Schizophrenia (BACS) and Positive

and Negative Syndrome Scale (PANSS) instruments to capture “core” PSD psychopathology dimensions (Fig. 1A-B). We observed differences across DSM diagnoses (Fig. 1A, $p < 0.05$, Bonferroni corrected), however, symptom distributions revealed notable overlap across the PSD sample that crossed diagnostic boundaries (11, 32). Furthermore, we observed marked collinearity between symptom items across the PSD sample (Fig. 1B), indicating that a lower-dimensional solution may better capture this symptom space. Specifically, we hypothesized that such a lower-dimensional symptom solution may improve PSD brain-behavior mapping as compared to a high-dimensional solution or pre-existing symptom scales. Here we report results from a principal component analysis (PCA) as it produces a deterministic solution with orthogonal axes (i.e. no *a priori* number of factors needs to be specified) and captures all symptom variance.

Results were highly consistent with prior symptom-reduction studies in PSD: we identified 5 PCs (Fig. 1C), which captured ~50.93% of all variance (see Methods & Fig. S2) (23). Notably, complementary data-reduction procedures produced similar effects for this PSD sample (e.g. independent component analysis (ICA), see Supplementary Note 1 & Fig. S3). The key innovation here is the combined analysis of canonical psychosis symptoms (i.e. positive and negative) and cognitive deficits, which are a fundamental PSD feature (17). The 5 PCs revealed few distinct boundaries between DSM categories (Fig. 1D). Fig. 1E highlights symptom configurations forming each PC, which drove their naming via the pattern of the most strongly weighted individual symptoms (e.g. PC3 – “Psychosis Configuration”; see Supplementary Note 2). Critically, PC axes were not parallel with traditional aggregate symptom scales. For instance, PC3 is angled at ~45° to the dominant direction of PANSS Positive and Negative symptom variation (purple and blue arrows respectively in Fig. 1F).

Next, we show that the PCA solution was highly stable when tested across sites, k-fold cross-validations, and split-half replications (see Supplementary Note 3 & Fig. S4-S5). Importantly, results were not driven by medication status or dosage (Fig. S7). Collectively, these data-reduction analyses strongly support a stable and replicable low-rank PSD symptom geometry.

Dimensionality-Reduced PSD Symptom Geometry Reveals Novel and Robust Neuro-Behavioral Relationships.

Next, we tested if the dimensionality-reduced symptom geometry can identify robust and novel patterns of neural variation across the PSD. Across all neural analyses we used global brain connectivity (GBC), the column-wise mean of the full FC matrix because it yields a parsimonious metric reflecting how “globally” coupled a particular area is to the rest of the brain (34) (see Methods). Furthermore, we selected GBC because: i) the metric is agnostic regarding the location of “dysconnectivity” as it weights each area equally; ii) it yields an interpretable dimensionality-reduction of the full FC matrix; iii) unlike the full FC matrix or other abstracted measures, GBC produces a neural map, which can be related

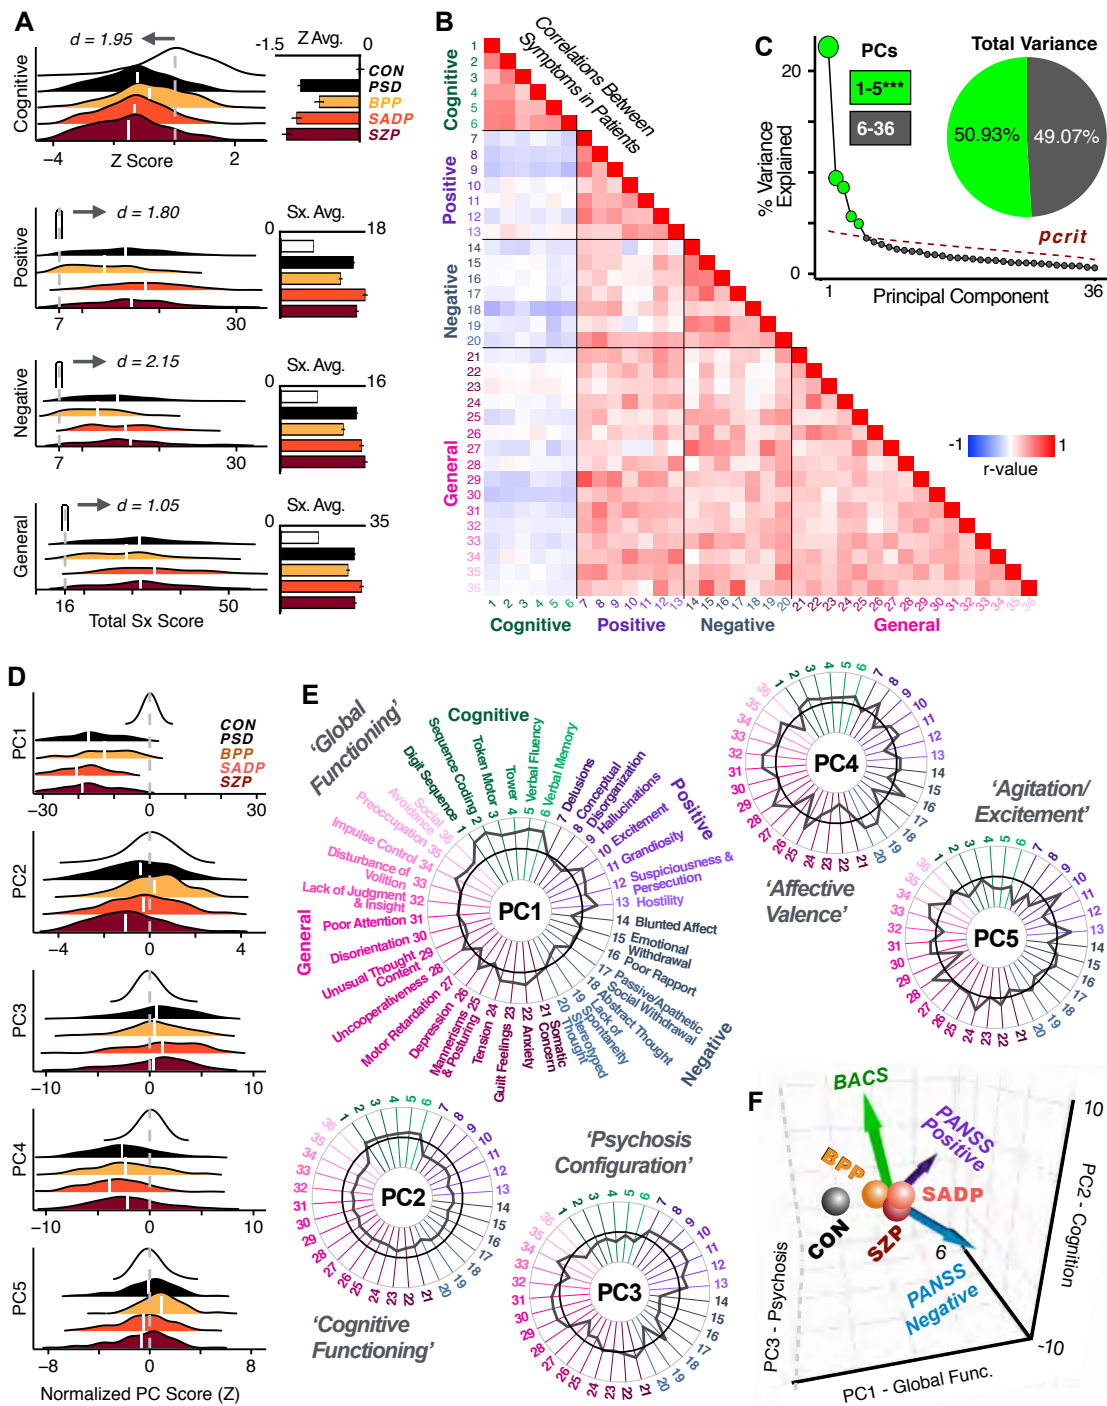


Fig. 1. Quantifying data-driven low-dimensional variation of cross-diagnostic psychosis spectrum disorder symptoms and cognitive deficits. (A) Distributions of symptom severity for each of the DSM diagnostic groups across core psychosis symptoms (PANSS positive, negative, and general symptoms tracking overall illness) and cognitive deficits (BACS composite cognitive performance). BPP: bipolar disorder with psychosis (yellow, N=150); SADP: schizo-affective disorder (orange, N=119); SZP - schizophrenia (red, N=167); All PSD patients (black, N=436); Controls (white, N=202). Bar plots show group means; error bars show standard deviations. (B) Correlations between 36 psychosis and cognitive deficit symptoms for all PSD patients (N=436). (C) Screeplot showing the % variance explained by each of the principal components (PCs) from a principal component analysis performed using all 36 symptom measures across all 436 PSD patients. The size of each point is proportional to the variance explained. The first five PCs (green) survived permutation testing ($p < 0.05$, 5,000 permutations). Together they capture 50.93% of all symptom variance (inset). (D) Distribution plots showing subject scores for the 5 significant PCs for each of the clinical groups, normalized relative to the control group. Note that control subjects (CON) were not used to derive the PCA solution. However, all subjects, including controls, can in principle be projected into the data-reduced symptom geometry. (E) Loading profiles for the 36 PANSS/BACS symptom measures on the 5 significant PCs. Each PC ("Global Dysfunction", "Cognition", "Psychosis Configuration", "Affective Valence", "Agitation/Excitement") was labeled according to the pattern of symptom loadings (see **Supplementary Note 2**). See **Fig. S2G** for numerical values of the PC loadings. (F) PCA solution projected into a 3D coordinate space defined by the first three significant PC axes. Colored arrows show *a priori* aggregate PANSS/BACS vectors projected into the data-reduced PC1-3 coordinate space. The *a priori* aggregate symptom vectors do not directly align with data-driven PC axes, highlighting that PSD symptom variation is not captured fully by any one aggregate *a priori* symptom score. Spheres denote centroids (i.e. center of mass) for each of the patient diagnostic groups and control subjects. Alternative views showing individual patients and controls projected into the PCA solution are shown in **Fig. S2A-F**.

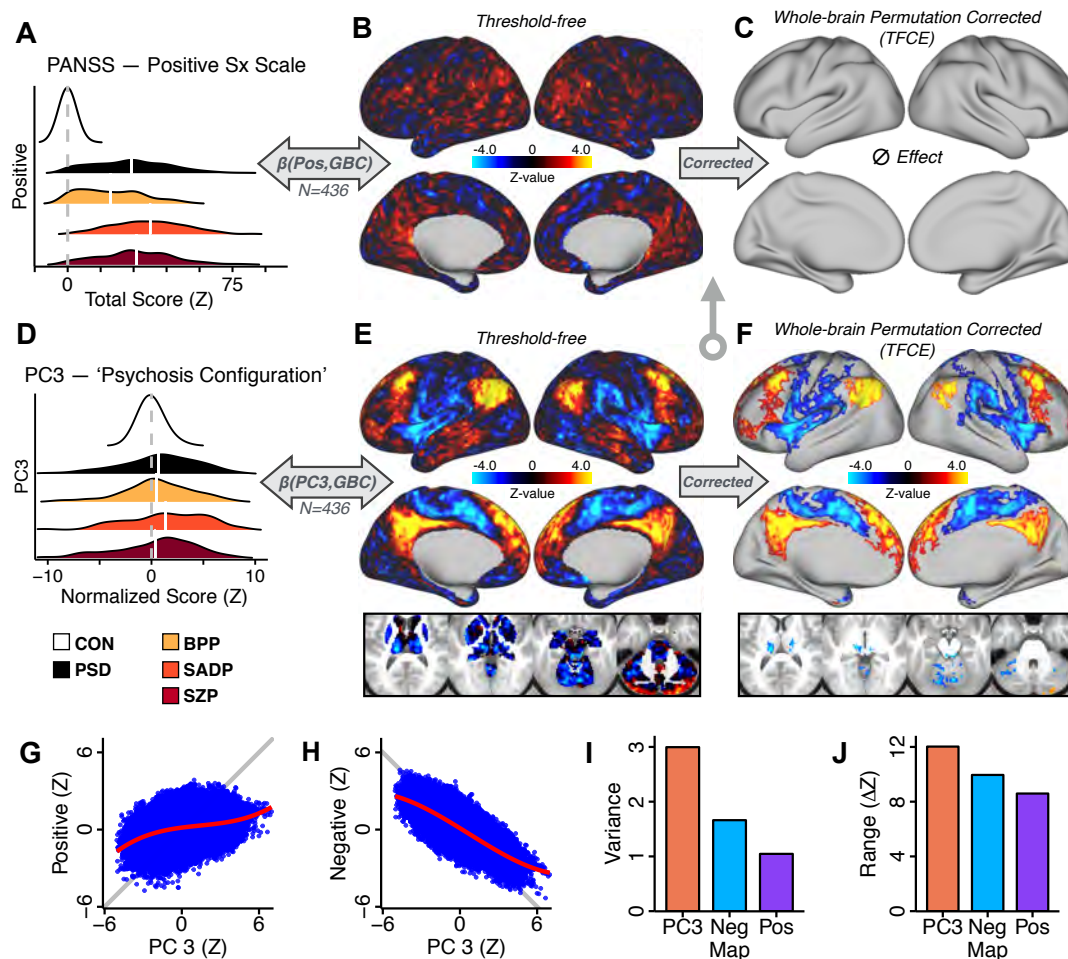


Fig. 2. Dimensionality-reduced symptom variation reveals robust neuro-behavioral mapping. (A) Distributions of total PANSS Positive symptoms for each of the clinical diagnostic groups normalized relative to the control group (white=CON; black=all PSD patients; yellow=BPP; orange=SADP; red=SZP). (B) $\beta_{\text{Positive GBC}}$ map showing the relationship between the aggregate PANSS Positive symptom score for each patient regressed onto global brain connectivity (GBC) across all patients (N=436). (C) No regions survived non-parametric family-wise error (FWE) correction at $p < 0.05$ using permutation testing with threshold-free cluster enhancement (TFCE). (D) Distributions of scores for PC3 "Psychosis Configuration" across clinical groups, again normalized to the control group. (E) $\beta_{\text{PC3 GBC}}$ map showing the relationship between the PC3 "Psychosis Configuration" score for each patient regressed onto GBC across all patients (N=436). (F) Regions surviving $p < 0.05$ FWE whole-brain correction via TFCE showed clear and robust effects. (G) Comparison between the Psychosis Configuration symptom score versus the aggregate PANSS Positive symptom score GBC map for every datapoint in the neural map (i.e. greyordinate in the CIFTI map). The sigmoidal pattern indicates an improvement in the Z-statistics for the Psychosis Configuration symptom score map (panel E) relative to the aggregate PANSS Positive symptom map (panel B). (H) A similar effect was observed when comparing the Psychosis Configuration GBC map relative to the PANSS Negative symptoms GBC map (Fig. S8). (I) Comparison of the variances for the Psychosis Configuration, PANSS Negative and PANSS Positive symptom map Z-scores. (J) Comparison of the ranges between the Psychosis Configuration, Negative and Positive symptom map Z-scores. Behavior-to-neural maps for all 5 PCs and all 4 traditional symptom scales (BACS and PANSS subscales) are shown in Fig. S8.

to other independent neural maps (e.g. gene expression or pharmacology maps, discussed below).

All 5 PCs captured unique GBC variation patterns across the PSD (Fig. S8). We highlight the PC3 solution (i.e. PC3 "Psychosis Configuration") to illustrate the benefit of the low-rank PSD symptom geometry for neuro-behavioral mapping relative to traditional aggregate PANSS symptom scales. The relationship between total PANSS Positive scores and GBC across N=436 PSD patients (Fig. 2A) was statistically modest (Fig. 2B) and surprisingly no areas survived whole-brain type-I error protection (Fig. 2C, $p < 0.05$). In contrast, regressing PC3 data-driven scores onto GBC across N=436 patients revealed a robust symptom-to-neural $\beta_{\text{PC3 GBC}}$ map (Fig. 2E-F), which survived whole-brain type-I error protection. Of note, the PC3 "Psychosis Configuration" axis is bi-directional whereby individuals who score either highly

positively or negatively are symptomatic. Therefore, a high positive PC3 score was associated with both reduced GBC across insular and superior dorsal cingulate cortices, thalamus, and anterior cerebellum and elevated GBC across precuneus, medial prefrontal, inferior parietal, superior temporal cortices and posterior lateral cerebellum – consistent with the "default-mode network" (35). A high negative PC3 score would exhibit the opposite pattern. Critically, this robust symptom-to-neural mapping emerged despite no DSM diagnostic group differences in PC3 scores (Fig. 2D). The diverging nature of this axis may be captured in the ICA solution, when orthogonality is not enforced (Fig. S11).

Moreover, the PC3 neuro-behavioral map exhibited improved statistical properties relative to other GBC maps computed from traditional aggregate PANSS symptom scales (Fig. 2G-J).

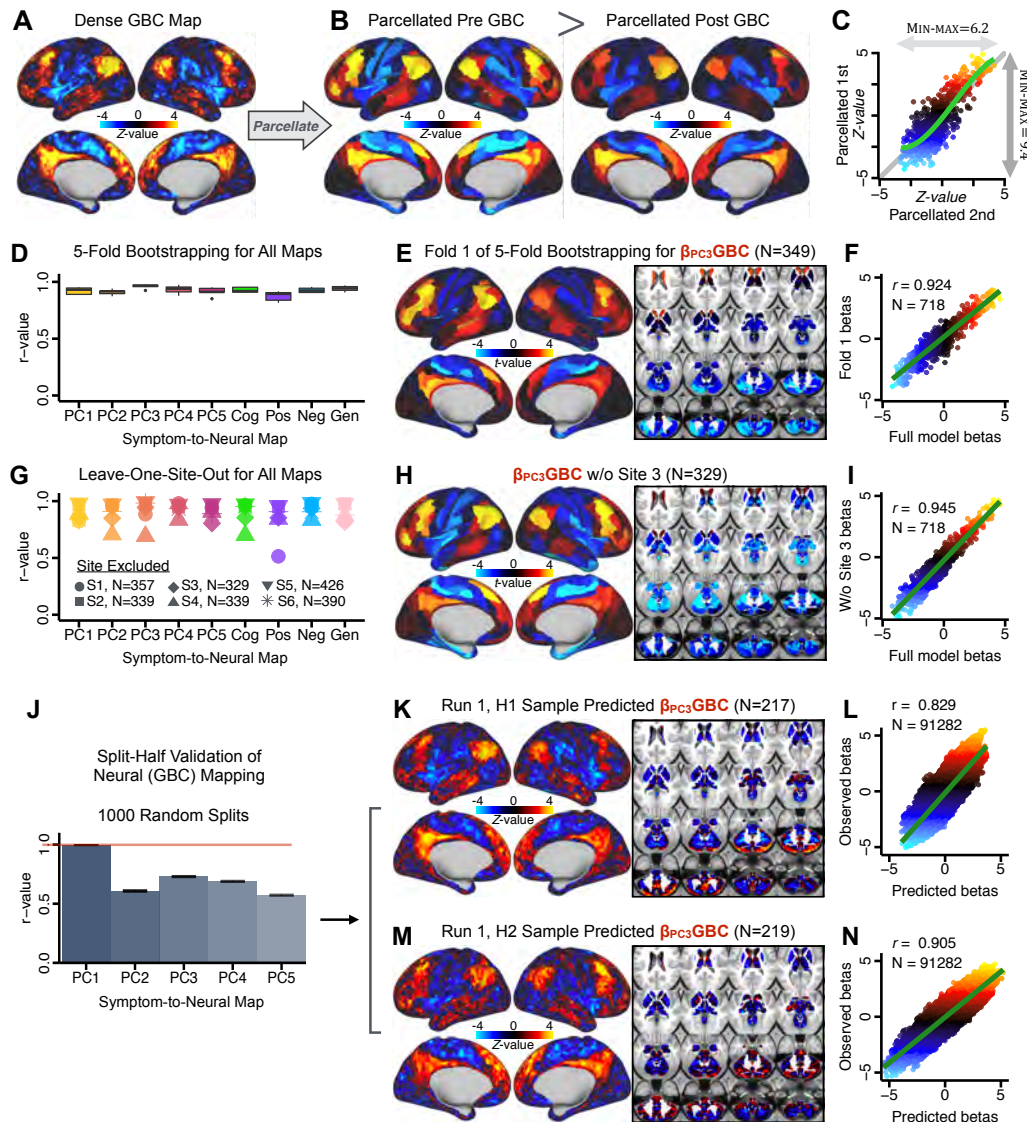


Fig. 3. Parcellated neuro-behavioral GBC maps reflecting psychosis configuration are statistically robust and replicable. (A) Z-scored *PC3 Psychosis Configuration* GBC neural map at the 'dense' (full CIFTI resolution) level. (B) Neural data were parcellated using a whole-brain functional parcellation (33) before computing GBC at the patient level, which yields stronger statistical values in the Z-scored *Psychosis Configuration* GBC neural map as compared to when parcellation was performed after computing GBC for each patient. (C) A comparison of pre- and post-GBC parcellation *Psychosis Configuration* GBC maps shows a statistical improvement if parcellation is applied prior to computing GBC at the single patient level. This effects indicates neural feature data reduction via a functional neural parcellation (33) improves signal-to-noise for the group-level statistical map. (D) Summary of similarity for all symptom-to-GBC maps (PCs and traditional symptom scales) across 5-fold cross-validation. Boxplots show the range of r values between coefficient maps for each fold and the full model. (E) Normalized coefficient map from regression of individual patients' *Psychosis Configuration* scores onto parcellated GBC data, shown here for a subset of patients from Fold 1 out of 5 (N=349). The greater the magnitude of the coefficient for a parcel, the stronger the statistical relationship between GBC of that parcel and *Psychosis Configuration* PC score. (F) Correlation between the value of each parcel in the regression model computed using patients in Fold 1 and the full PSD sample (N=436) model. Each point is a single parcel (718 parcels total). The resulting normalized coefficient map from the leave-one-fold-out model was highly similar to the map obtained from the full PSD sample model ($r=0.924$). (G) Summary of leave-one-site-out regression for all symptom-to-GBC maps. Regression of symptom scores onto parcellated GBC data, each time leaving out subjects from one site, resulted in highly similar maps. This highlights that the relationship between *Psychosis Configuration* and GBC variation is robust and not driven by a specific site. (H) Normalized coefficient map from regression of individual subject scores along *Psychosis Configuration* onto neural parcellated GBC data for all patients except one site. For example, Site 3 is excluded here given that it recruited the most patients (and therefore may have the greatest statistical impact on the full model). (I) Correlation between the value of each parcel in the regression model computed using all patients minus Site 3, and the full PSD sample model, again indicating a consistent neuro-behavioral relationship. (J) Split-half replication of *Psychosis Configuration*-to-GBC mapping. For each stratified split-half replication run, the full sample was randomly split into two halves (H1 and H2) with the proportion of each diagnostic group (BPP, SADP, SZP) preserved within each half. PCA scores for subjects in each half are computed using the PC loadings from the PCA conducted using subjects in the other half. Observed PCA scores are computed from a PCA on the same half-sample of subjects. Each set of scores were then regressed against parcellated GBC for subjects in a given half, resulting in a coefficient map for each PC (as in Fig. 2) reflecting the strength of the relationship between PC score and GBC across subjects. The coefficient maps for predicted and observed scores were then cross-correlated for PCs 1-5. This process was repeated 1,000 times. Bar plots show the mean correlation across the 1,000 runs; error bars show standard error of the mean. Note that the split-half effect for PC1 was exceptionally robust. The split-half consistency for PC3, while lower, was still highly robust and well above chance. (K) Normalized coefficient map from *Psychosis Configuration*-to-GBC regression for the first half (H1) patients, shown here for one exemplar run out of 1000 split-half validations. (L) Correlation across 718 parcels between the H1 predicted coefficient map (i.e. panel K) and the observed coefficient map for H1 patient sample. (M-N) The same analysis as K-L is shown for patients in H2, indicating a striking consistency in *Psychosis Configuration* neural map correspondence.

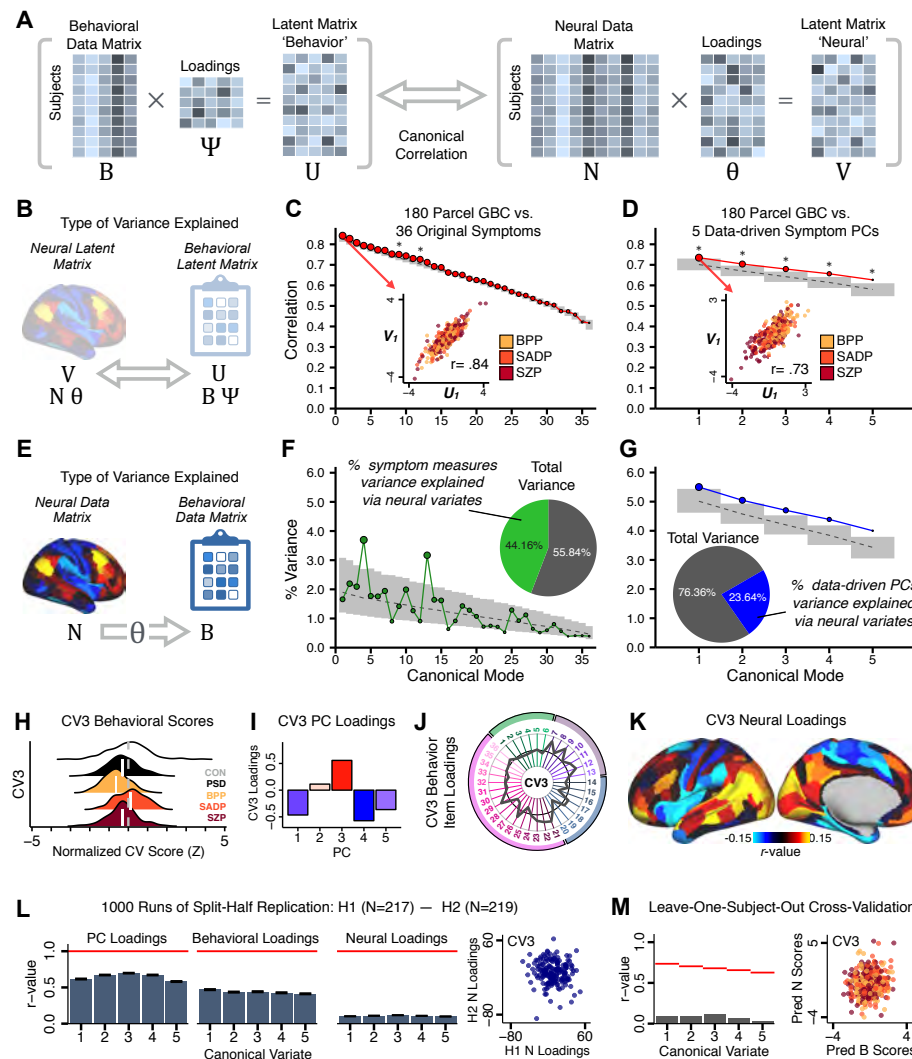


Fig. 4. Multivariate behavioral-to-neural feature mapping using canonical correlation analysis (CCA). (A) CCA computes a relationship between two sets of observed variables (in this case, a matrix **B** of behavioral measures and a matrix **N** of neural measures across all patients), by solving for two transformation matrices Ψ and Θ such that the correlations between the latent variables are maximized (represented as columns along the transformed 'latent' matrices **U** and **V**). Here each column in **U** and **V** is referred to as a canonical variate (CV); each corresponding pair of CVs (e.g. U_1 and V_1) is referred to as a canonical mode. (B) CCA maximized correlations between the CVs (i.e. matrices **U** and **V**) (C) Screeplot showing canonical modes for the CVs obtained from 180 neural features (cortical GBC symmetrized across hemispheres) and 36 single-item PANSS and BACS symptom measures. Inset illustrates the correlation ($r=0.85$) between the CV of the first mode, U_1 and V_1 (note that the correlation was not driven by a separation between diagnoses). Correlations between each subsequent pair of variates is computed from the residuals of the previous pair. Modes 9 and 12 remained significant after FDR correction across all 36 modes. (D) CCA was obtained from 180 neural features and 5 low-dimensional symptom scores derived via the PCA analysis. Here all modes remained significant after FDR correction. Dashed black line shows the null calculated via a permutation test with 5,000 shuffles; grey bars show 95% confidence interval. (E) Correlation between the behavioral data matrix **B** and the neural data matrix weighted by the transformation matrix ($N\Theta$) reflects the amount of variance in **B** that can be explained by the final latent neural matrix **V**. Put differently, this transformation calculates how much of the symptom variation can be explained by the latent neural features. (F) Proportion of symptom variance explained by each of the neural CVs in a CCA performed between 180 neural features and all 36 behavioral measures. Inset shows the total proportion of behavioral variance explained by the neural variates. (G) Proportion of total behavioral variance explained by each of the neural CVs in a CCA performed between 180 neural features and the 5 low-dimensional symptom scores derived via the PCA analysis. While CCA using symptom PCs has fewer dimensions and thus lower total variance explained (see inset), each neural variate explains a higher amount of symptom variance than seen in F, suggesting that CCA could be further optimized by first obtained a principled low-rank symptom solution. Dashed black line shows the null calculated via a permutation test with 5,000 shuffles; grey bars show 95% confidence interval. Behavioral-to-neural mapping is shown in Fig. S14. (H) Characterizing CV symptom configurations using CV3 as an example. Distributions of CV3 scores by each DSM diagnostic group. All scores are normalized to controls. (I) Behavioral canonical factor loadings for CV3. (J) Loadings of the original symptom items for CV3. (K) Neural canonical factor loadings for CV3. (L) Within-sample CCA cross-validation appeared robust (see Fig. S16). However, a split-half CCA replication using two independent non-overlapping patient samples was not reliable. Bar plots show the mean correlation for each CV between the first half (H1) and the second half (H2) CCA, each computed across 1,000 runs. Left: split-half replication of the behavioral PC loadings matrix Ψ ; Middle: individual behavioral item loadings; and Right: the neural loadings matrix Θ . Error bars show the standard error of the mean. The neural loadings matrix Θ in particular was not stable. Scatterplot shows the correlation between CV3 neural loadings for H1 vs. H2 for one example CCA run, illustrating lack of reliability. (M) Leave-one-subject-out cross-validation further highlights CCA instability. Here, CCA was computed for all except one patient ($N=435$). The loadings were then used to compute the predicted neural (**N**) and behavioral (**B**) scores for the left-out patient. This was repeated across all patients. This yielded predicted neural and predicted behavioral scores for each patient for each of the five CVs. These correlations were far lower than the canonical correlation values obtained in the full CCA model (shown as red lines). This is highlighted in the scatterplot at far right for CV3 ($r=0.12$, whereas the full sample CCA was $r=0.68$). CCA solutions for subcortex parcels only (Fig. S12) and network-level neural feature reduction Fig. S13 are shown in the Supplement.

Univariate Neuro-Behavioral Map of Psychosis Configuration is Reproducible. After observing improved PC3 neuro-behavioral statistics, we tested if these symptom-to-neural map patterns replicate. Recent attempts to derive stable symptom-to-neural mapping using multivariate techniques, while inferentially valid, have not replicated (25). This fundamentally relates to the tradeoff between the sample size needed to resolve multivariate neuro-behavioral solutions and the size of the feature space. To mitigate the feature size issue we re-computed the β_{PCGBC} maps using a functionally-derived whole-brain parcellation via the recently-validated CAB-NP atlas (33, 36) (Methods). Here, a functional parcellation is a principled way of neural feature reduction (to 718 parcels) that can also appreciably boost signal-to-noise (33, 36). Indeed, parcellating the full-resolution “dense” resting-state signal for each subject prior to computing GBC statistically improved the group-level neuro-behavioral maps compared to parcellating after computing GBC (Fig. 3A–C, all maps in Fig. S9). Results demonstrate that the univariate symptom-to-neural mapping was indeed stable across 5-fold bootstrapping, leave-site-out, and split-half cross-validations (Fig. 3D–N, see **Supplementary Note 4**), yielding consistent neuro-behavioral PC3 maps. Importantly, the neuro-behavioral maps computed using ICA showed comparable results (Fig. S11).

A Multivariate PSD Neuro-Behavioral Solution Can be Computed but is Not Reproducible with the Present Sample Size. Several studies have reported “latent” neuro-behavioral relationships using multivariate statistics (15, 37, 38), which would be preferable because they maximize covariation across neural and behavioral features simultaneously. However, concerns have emerged whether such multivariate results will replicate due to the size of the feature space (25). Nevertheless, we tested if results improve with canonical correlation analysis (CCA) (39) which maximizes relationships between linear combinations of symptom (*B*) and neural measures (*N*) across all subjects (Fig. 4A).

We examined two CCA solutions using symptom scores in relation to neural features: i) all item-level scores from the PANSS/BACS (Fig. 4C&F); ii) data-reduced PC scores (Fig. 4D&G, see Methods). In turn, to evaluate if the number of neural features affects the solution we computed CCA using: i) 718 bilateral whole-brain parcels derived from the brain-wide cortico-subcortical CAB-NP parcellation (33, 36); ii) 359 bilateral subcortex-only parcels; iii) 192 symmetrized subcortical parcels; iv) 180 symmetrized cortical parcels; v) 12 functional networks defined from the parcellation (33, 36). Notably, 4 out of the 5 CCA solutions were not robust: 718 whole-brain & 359 subcortex-only parcel solutions did not produce stable results according to any criterion (Fig. S14), whereas the symmetrized subcortical (192 features, Fig. S12) and network-level (12 features, Fig. S13) solutions captured statistically modest effects relative to the 180 symmetrized cortical CCA results (Fig. 4B–D). Therefore, we characterized the 180-parcel CCA solution further. Only 2 out of 36 CCA modes for the item-level symptom solution survived permutation testing and false discovery rate (FDR) correction

(Fig. 4C). In contrast, all 5 modes computed using PC scores survived both permutation and FDR correction (Fig. 4D). Critically, we found that no single CCA mode computed on item-level symptoms captured more variance than the CCA modes derived from PC scores, suggesting that the PCA-derived dimensions capture more neurally-relevant variation than any one single clinical item (Fig. 4E–G). Additional detailed CCA results are presented in **see Supplementary Note 5** and **Fig. S14**.

We highlight an example canonical variate (CV) solution across both behavioral and neural effects for CV3. We show the CV3 scores across diagnostic groups normalized to controls Fig. 4H as well as how the CV3 loads onto each PC (Fig. 4I, full results are shown in Fig. S15). The negative loadings on to PCs 1, 4, and 5 and the high positive loadings on to PC3 in Fig. 4I indicate that CV3 captures some shared variation across PCs. This can also be visualized by computing how the CV3 projects onto the original 36 behavioral items (Fig. 4J). Finally, the neural loadings for CV3 are shown in Fig. 4K and data for all 5 CVs are shown in Fig. S14.

Next, we tested if the 180-parcel CCA solution is stable and replicable, as done with PC-to-GBC univariate results. The CCA solution was robust when tested with k-fold and leave-site-out cross-validation (Fig. S16) because these methods use CCA loadings derived from the full sample. However, the CCA loadings did not replicate in non-overlapping split-half samples (Fig. 4L, see **see Supplementary Note 6**). Moreover, a leave-one-subject-out cross-validation revealed that removing a single subject from the sample affected the CCA solution such that it did not generalize to the left-out subject (Fig. 4M). As noted, the PCA-to-GBC univariate mapping was substantially more reproducible for all attempted cross-validations relative to the CCA approach, reflecting the fact that substantially more power is needed to resolve a stable multi-variate neuro-behavioral effect with this many features (40). Therefore, we leverage the univariate neuro-behavioral result for subsequent subject-specific model optimization and comparisons to molecular neuroimaging maps.

A Major Proportion of Overall Neural Variance May Not be Relevant for Psychosis Symptoms. Most studies look for differences between clinical and control groups, but to our knowledge no study has tested whether both PSD and healthy controls actually share a major portion of neural variance that may be present across all people. If the bulk of the neural variance is similar across both PSD and CON groups then including this clinically-irrelevant neural signal might obscure neuro-behavioral relationships that are clinically relevant. To test this, we examined the shared variance structure of the neural signal for all PSD patients (N=436) and all controls (N=202) independently by conducting a PCA on the GBC maps (see Methods). Patients’ and controls’ neural signals were highly similar for each of the first 3 neural PCs (>30% of all neural variance in each group) (Fig. S17A–J). These PCs may reflect a “core” shared symptom-irrelevant neural variance that generalizes across all people. These data suggest that the bulk of neural variance in PSD patients may

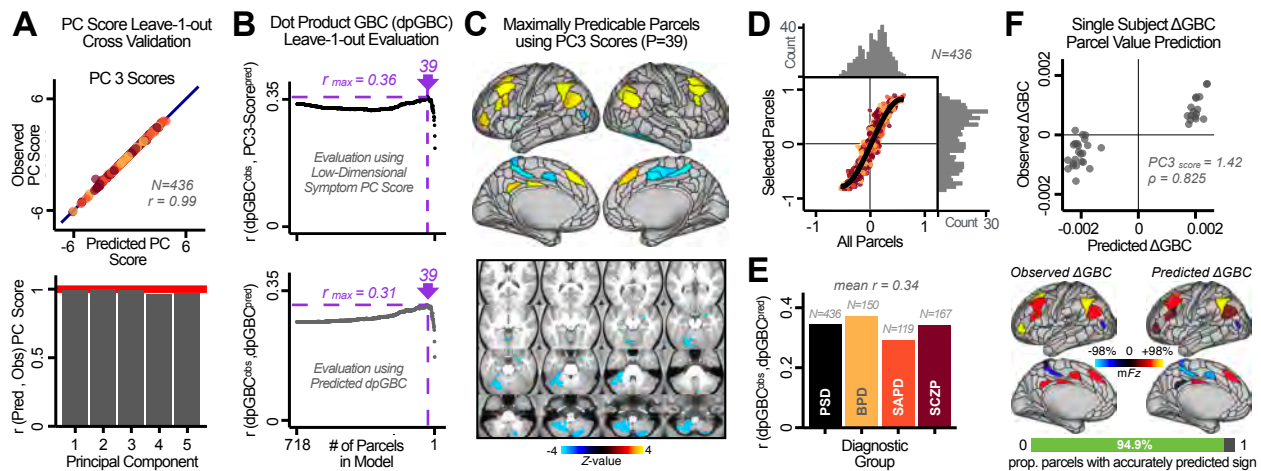


Fig. 5. Optimizing neural feature selection to inform single-subject prediction via a low-dimensional symptom solution. (A) Leave-one-out cross-validation for the symptom PCA analyses indicates robust individual score prediction. Top panel: Scatterplot shows the correlation between each subject's predicted $PC3$ score from a leave-one-out PCA model and their observed $PC3$ score from the full-sample PCA model, $r=0.99$. Bottom panel: Correlation between predicted and observed individual PC scores was above 0.99 for each of the significant PCs (see Fig. 1). The red line indicates $r=1$. (B) We developed a univariate step-down feature selection framework to obtain the most predictive parcels using a subject-specific approach via the $dpGBC$ index. Specifically, the 'observed' patient-specific $dpGBC^{obs}$ was calculated using each patient's ΔGBC^{obs} (i.e. the patient-specific GBC map vs. the group mean GBC for each parcel) and the 'reference' symptom-to-GBC $PC3$ map (described in Fig. 3B) [$dpGBC^{obs} = \Delta GBC^{obs} \cdot \beta_{PC3} GBC^{obs}$]. See Methods & Fig. S18 for complete feature selection procedure details. In turn, we computed the predicted $dpGBC$ index for each patient by holding their data out of the model and predicting their score ($dpGBC^{pred}$). We used two metrics to evaluate the maximally predictive feature subset: i) The correlation between $PC3$ symptom score and $dpGBC^{obs}$ across all $N=436$, which was maximal for $P = 39$ parcels [$r=0.36$, purple arrow]; ii) The correlation between $dpGBC^{obs}$ and $dpGBC^{pred}$, which also peaked at $P = 39$ parcels [$r=0.31$, purple arrow]. (C) The $P = 39$ maximally predictive parcels from the $\beta_{PC3} GBC^{obs}$ map are highlighted (referred to as the "selected" map). (D) Across all $n=436$ patients we evaluated if the selected parcels improve the statistical range of similarities between the ΔGBC^{obs} and the $\beta_{PC3} GBC^{obs}$ reference for each patient. For each subject the value on the X-axis reflects a correlation between their ΔGBC^{obs} map and the $\beta_{PC3} GBC^{obs}$ map across all 718 parcels; the Y-axis reflects a correlation between their ΔGBC^{obs} map and the $\beta_{PC3} GBC^{obs}$ map only within the "selected" 39 parcels. The marginal histograms show the distribution of these values across subjects. (E) Each DSM diagnostic group showed comparable correlations between predicted and observed $dpGBC$ values. (F) Scatterplot for a single patient with a positive behavioral loading ($PC3$ score = 1.42) and also with a high correlation between predicted ΔGBC^{pred} versus observed ΔGBC^{obs} values for the "selected" 39 parcels ($p = 0.825$). Right panel highlights the observed vs. predicted ΔGBC map for this patient, indicating that 94.9% of the parcels were predicted in the correct direction (i.e. in the correct quadrant).

actually not be symptom-relevant, which highlights the importance of optimizing symptom-to-neural mapping.

Optimizing Neuro-Behavioral Mapping Features for Personalized Prediction via Dimensionality-reduced Symptoms.

Above we demonstrate that PC scores can be reliably predicted across sites and cross-validation approaches (Fig. S4). Here we show that leave-one-subject-out cross-validation yields reliable effects for the low-rank symptom PCA solution (Fig. 5A). This stable single-subject PC score prediction provides the basis for testing if the derived neural maps can yield an individually-reliable subset of features. To this end, we developed a univariate quantitative neural feature selection framework (Fig. S18) based on data-reduced PC scores (i.e. $PC3$ score). Specifically, we computed a dot product GBC metric ($dpGBC$) that provides an index of similarity between an individual ΔGBC topography relative to a "reference" group-level PC GBC map (see Methods and Fig. S18). Using this $dpGBC$ index we found, via a feature selection step-down regression, a subset of parcels for which the symptom-to-neural feature statistical association was maximal (Fig. 5A). For $PC3$ we found $P = 39$ maximally predictive parcels out of the group map. Specifically, the relationship between $PC3$ symptom scores and $dpGBC$ values across subjects was maximal (Fig. 5B, top panel, $r = 0.36$) as was the relationship between predicted $dpGBC^{pred}$ vs. observed $dpGBC^{obs}$ (Fig. 5B, bottom

panel, $r = 0.31$) (see Fig. S19 for all PCs). Importantly, the "subset" feature map (i.e. $[\beta_{PC3} GBC^{obs}]^{P=39}$, Fig. 5C) exhibited improved statistical properties relative to the full map (i.e. $[\beta_{PC3} GBC^{obs}]^{P=718}$, Fig. 5D). Furthermore, the relationship between observed vs. predicted subset feature maps (i.e. $r[dpGBC^{obs}, dpGBC^{pred}]$) was highly consistent across DSM diagnoses (Fig. 5E). Finally, a single patient is shown for whom the correlation between their predicted and observed subset feature maps was high (i.e. $r[\Delta GBC^{pred}, \Delta GBC^{obs}]$, Fig. 5F), demonstrating that the dimensionality-reduced symptom scores can be used to quantitatively optimize individual neuro-behavioral map features.

Single Patient Evaluation via Neuro-Behavioral Target Map Similarity.

We demonstrated a quantitative framework for computing neuro-behavioral model at the single-patient level. This brain-behavior space (BBS) model was optimized along a single dimensionality-reduced symptom axis (i.e. PC score). Next, we tested a hybrid patient selection strategy by first imposing a PC-based symptom threshold, followed by a target neural similarity threshold driven by the most highly predictive neuro-behavioral map features. Specifically, the "neural similarity prediction index (NSPI)" computes a patient-specific Spearman's ρ between that patient's ΔGBC^{obs} and the group reference $\beta_{PC3} GBC^{obs}$ map using the maximally predictive $P = 39$ parcels (see Fig. S20 for whole-brain results and alternative similarity metrics). Fig.

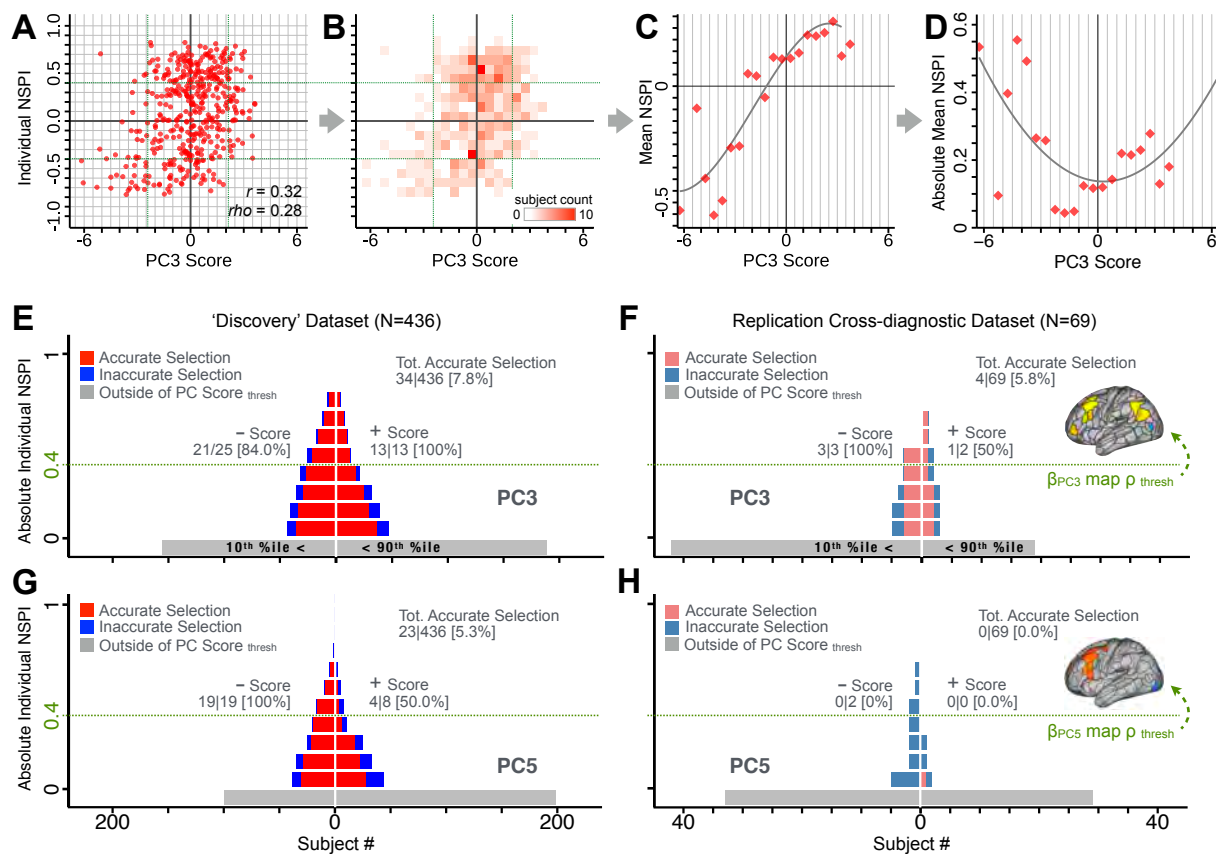


Fig. 6. Evaluating patient-specific similarity to derived neuro-behavioral targets via data-reduced symptom scores. (A) Scatterplot of $PC3$ symptom score (X-axis) versus $PC3$ neural similarity prediction index (NSPI, Y-axis) for all 436 PSD subjects. The NSPI is defined as the Spearman's correlation between ΔGBC^{obs} and the $\beta_{PC3} GBC^{obs}$ map of the maximally-predictable "selected" $P = 39$ parcels. Alternative metrics are shown in Fig. S20. (B) Bins across axes express subject counts within each cell as a heatmap, indicating a high similarity between symptom PC score and $PC3$ NSPI for a number of patients. (C) Mean NSPI is computed for a given bin along the X-axis to visualize patient clustering. Note the sigmoidal shape of the distribution reflecting greater neural similarity at more extreme values of the $PC3$ score. (D) The absolute value of the mean NSPI reflects the magnitude irrespective of neural similarity direction. This highlights a quadratic effect, showing that patients with higher $PC3$ symptom scores (either positive/negative) exhibited higher neural correspondence of their maps with the target neural reference map. (E) Using the NSPI and PC scores we demonstrate one possible brain-behavioral patient selection strategy. We first imposed a PC score symptom threshold to select patients at the extreme tails (i.e. outside of the 10th – 90th percentile behavioral range [$>+2.17$ or <-2.41]). Note that this patient selection strategy excludes patients (shown in grey) below the PC symptom score threshold. This yielded $n=38$ patients. Next, for each patient we predicted the sign of their individual NSPI based on their individual $PC3$ score, which served as the basis for the neural selection. Next, at each NSPI threshold we evaluated the proportion of patients correctly selected until there were no inaccurately selected patients in at least one $PC3$ direction (green line or higher). The number of accurately (A) vs. inaccurately (I) selected patients within each bin is shown in red and blue respectively. Note that as the neural ρ threshold increases the A/I ratio increases. (F) The neural and behavioral thresholds defined in the "discovery" sample were applied to an independent "replication" dataset ($N=69$, see Methods), yielding a similar final proportion of accurately selected patients. (G) The same brain-behavioral patient selection strategy was repeated for $PC5$ in the discovery sample (thresholds of 10th percentile = -1.89 and 90th percentile = +1.47; NSPI threshold of $\rho=0.4$ optimized for $PC5$). Results yielded similar A/I ratios as found for $PC3$. (H) The neural and behavioral thresholds for $PC5$ defined in the discovery sample were applied to the replication sample. Here the results failed to generalize due to true clinical differences between the discovery and replication samples.

6A shows a significant relationship between each patient's $PC3$ symptom score (X-axis) and the neural similarity index (Y-axis). In turn, **Fig. 6B** shows binned results, which provides a visual intuition for patient segmentation across both the neural and behavioral indices (**Fig. 6A**, right: binned by $\rho = 0.1$ & $PC3_{score} = 0.5$). For patients at either tail, the neuro-behavioral relationship was robust. Conversely, patients with a low absolute $PC3$ score showed a weak relationship with symptom-relevant neural features. This is intuitive because these individuals do not vary along the selected PC symptom axis.

Fig. 6C shows the mean NSPI across subjects within each PC symptom bin along the X-axis. The resulting sigmoid captures that patients exhibit greater neural similarity if their PC symptom scores are more extreme. To evaluate if this relationship can yield a personalized patient selection, we com-

puted the absolute NSPI **Fig. 6D**). The effect was approximated by a quadratic function, highlighting that patients with extreme $PC3$ scores (either positive or negative) exhibited a stronger NSPI (i.e. personalized neural effects that strongly resembled the reference ($\beta_{PC3} GBC^{obs}$)^P = 39 map). **Fig. 6E** shows the application of this neuro-behavioral selection procedure, demonstrating that PSD patients with extreme $PC3$ scores (defined at the top/bottom 10th percentile of the "discovery" sample, $+2.17 < PC3_{score} < -2.41$) exhibit high NSPI values. We observed an inherent trade-off such that if the PC score threshold was raised then neural target similarity confidence goes up, but fewer patients will be selected. In the discovery PSD sample, all patients were accurately selected above the following neuro-behavioral thresholds: $90^{th} \%tile < PC3_{score} < 10^{th} \%tile$ and $|\rho| > 0.4$ (**Fig. 6E**, 34/436 patients selected, green line). We show

consistent effects for when the selection was applied to $PC5$ (Fig. 6F; results for all PCs in Fig. S21).

To test if the neuro-behavioral selection is generalizable, we used an independent cross-diagnostic sample of 30 patients diagnosed with SZP and 39 diagnosed with obsessive-compulsive disorder (OCD) (Fig. 6G, see Methods and Table S3 for details). Applying the “discovery” selection thresholds yielded similar results for ~6% of the cross-diagnostic “replication” sample for $PC3$ (Fig. 6G, full analyses in Fig. S20C). Notably, no replication sample patients were selected along the neuro-behavioral thresholds for $PC5$ (Fig. 6H). While there are SZP patients in the replication cross-diagnostic sample, few scored highly on $PC5$ and none met the neural similarity threshold, emphasizing that not all patients within the same DSM-based diagnosis will exhibit variation along the same neuro-behavioral axis. Collectively, these results show that data-driven symptom scores can pinpoint individual patients for whom their neural variation strongly maps onto a target neural reference map. These data also highlight that both symptom and neural information for an independent patient can be quantified in the reference ‘discovery’ BBS using their symptom data alone.

Subject-Specific PSD Neuro-Behavioral Features Track Neuropharmacological Map Patterns.

Next, we use the personalized BBS selection in a proof-of-concept framework for informing molecular mechanism of possible treatment response by relating subject-specific BBS to independently-acquired pharmacological neuroimaging maps. Here we examine two mechanisms implicated in PSD neuropathology via ketamine, a N-methyl-D-aspartate (NMDA) receptor antagonist (42), and lysergic acid diethylamide (LSD), primarily a serotonin receptor agonist (28, 43, 44). We first quantified individual subjects’ BBS “locations” in the established reference neuro-behavioral geometry. The radarplot in Fig. 7A shows original symptoms whereas Fig. 7B shows ΔGBC^{obs} maps for two patients from the replication dataset (denoted here with X_{PC3} and Y_{PC3} , see Fig. S10 for other example patients). Both of these patients exceeded the neural and behavioral BBS selection indices for $PC3$ (defined independently in the “discovery” dataset, Fig. 6C). Furthermore, both patients exhibited neuro-behavioral variation in line with their expected locations in the BBS geometry. Specifically, patient X_{PC3} from the replication dataset scored highly negatively on the $PC3$ axis defined in the “discovery” PSD sample (Fig. 7A). In contrast, patient Y_{PC3} scored positively on the $PC3$ axis. Importantly, the correlation between the ΔGBC^{obs} map for each patient and the group-reference $\beta_{PC3}GBC$ was directionally consistent with their symptom PC score (Fig. 7B-C).

We then tested if the single-subject BBS selection could be quantified with respect to a neural map reflecting glutamate receptor manipulation, a hypothesized mechanism underlying PSD symptoms (45). Specifically, we used an independently collected ketamine infusion dataset, collected in healthy adult volunteers during resting-state fMRI (46). As with the clinical data, here we computed a ΔGBC map re-

flecting the effect of ketamine on GBC relative to placebo (Methods). The maximally-predictive $PC3$ parcels exhibited high spatial similarity with the ketamine map ($\rho=0.76$, see Methods), indicating that the ΔGBC pattern induced by ketamine tracks with the GBC pattern reflecting $PC3$ symptom variation.

Critically, because X_{PC3} is negatively loaded on the $PC3$ symptom axis, an NMDA receptor antagonist like ketamine may modulate symptom-relevant circuits in a way that reduces similarity with the $PC3$ map. This may in turn have an impact on the $PC3$ -like symptoms. Consistent with this hypothesis, X_{PC3} expresses predominantly depressive symptoms (Fig. 7A), and ketamine has been shown to act as an anti-depressant (41). This approach can be applied for patients that load along another axis, such as $PC5$. Fig. 7D-E shows the symptom and neural data for two patients whom met thresholds for $PC5$ selection (Fig. 6C). Notably, the selected $PC5$ map is anti-correlated with a ΔGBC map reflecting LSD vs. placebo effects (28) ($\rho=-0.44$, Fig. 7F). Hence areas modulated by LSD may map onto behavioral variation along $PC5$. Consequently, serotonergic modulation may be effective for treating Q_{PC5} and Z_{PC5} , via an antagonist or an agonist respectively. These differential similarities between pharmacological response maps and BBS maps (Fig. S22) can be refined for quantitative patient segmentation.

Group-Level PSD Neuro-Behavioral Features Track Neural Gene Expression Patterns.

To further inform molecular mechanism for the derived BBS results, we compared results with patterns neural gene expression profiles derived from the Allen Human Brain Atlas (AHBA) (30, 31) (Fig. 8A, Methods). Specifically, we tested if BBS cortical topographies, which reflect stable symptom-to-neural mapping along PSD, covary with the expression of genes implicated in PSD neuropathology. We focus here on serotonin and GABA receptors as well as interneuron markers (SST, somatostatin; PVALB, parvalbumin). Fig. 8B shows the distribution of correlations between the $PC3$ map and the cortical expression patterns of 20,200 available AHBA genes (other PCs shown in Fig. S23). Seven genes of interest are highlighted, along with their cortical expression topographies and their similarity with the $PC3$ BBS map (Fig. 8C-E). This BBS-to-gene mapping can potentially reveal novel therapeutic molecular targets for neuro-behavioral variation. For example, the HTR1E gene, which encodes the serotonin 5-HT_{1E} receptor, is highly correlated with the $PC3$ BBS map. This could drive further development of novel selective ligands for this receptor, which are not currently available (47).

Discussion

We found a robust and replicable symptom-to-neural mapping across the psychosis spectrum that emerged from a low-dimensional symptom solution. Critically, this low-rank symptom solution was predictive of a neural circuit pattern, which reproduced at the single-subject level. In turn, we show that the derived PSD symptom-to-neural feature maps exhibit spatial correspondence with independent pharmaco-

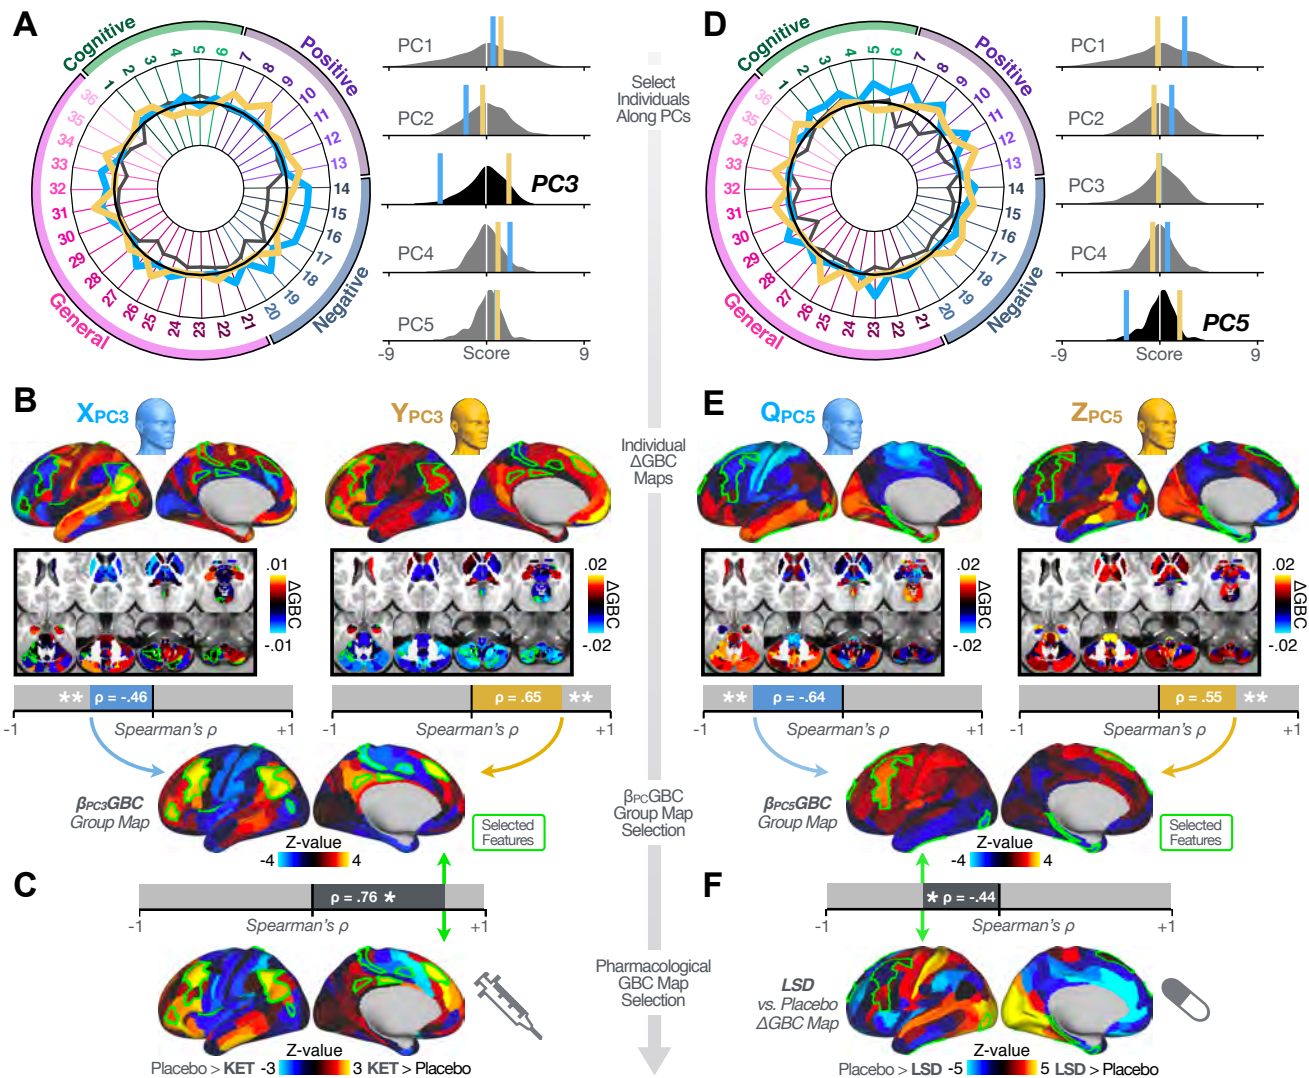


Fig. 7. Leveraging subject-specific brain-behavioral maps for molecular neuroimaging target selection. (A) Data for two individual patients from the replication dataset are highlighted for PC3: X_{PC3} (blue) and Y_{PC3} (yellow). Both of these patients scored above the neural and behavioral thresholds for PC3 defined in the “discovery” PSD dataset. Patient X_{PC3} loads highly negatively on the PC5 axis and Patient Y_{PC3} loads highly positively. Density plots show the projected PC scores for Patients X_{PC3} and Y_{PC3} overlaid on distributions of PC scores from the discovery PSD sample. (B) Neural maps show cortical and subcortical ΔGBC^{obs} for the two patients X_{PC3} and Y_{PC3} specifically reflecting a difference from the mean PC3. The similarity of ΔGBC^{obs} and the $\beta_{PC3}GBC^{obs}$ map within the most predictive neural parcels for PC3 (outlined in green). Note that the sign of neural similarity to the reference PC3 map and the sign of the PC3 score is consistent for these two patients. (C) The selected PC3 map (parcels outlined in green) is spatially correlated to the neural map reflecting the change in GBC after ketamine administration ($\rho = 0.76$, **Methods**). Note that Patient X_{PC3} who exhibits ΔGBC^{obs} that is anti-correlated to the ketamine map also expresses depressive moods symptoms (panel A). This is consistent with the possibility that this person may clinically benefit from ketamine administration, which may elevate connectivity in areas where they show reductions (41). In contrast, Patient Y_{PC3} may exhibit an exacerbation of their psychosis symptoms given that their ΔGBC^{obs} is positively correlation with the ketamine map. (D) Data for two individual patients from the discovery dataset are highlighted for PC5: Q_{PC5} (blue) and Z_{PC5} (yellow). Note that no patients in the replication dataset were selected for PC5 so both of these patients were selected from “discovery” PSD dataset for illustrative purposes. Patient Q_{PC5} loads highly negatively on the PC5 axis and Patient Z_{PC5} loads highly positively. Density plots show the projected PC scores for Patients Q_{PC5} and Z_{PC5} overlaid on distributions of PC scores from the discovery PSD sample. (E) Neural maps show cortical and subcortical ΔGBC^{obs} for Patients Q_{PC5} and Z_{PC5} , which are highly negatively and positively correlated with the selected PC5 map respectively. (F) The selected PC5 map (parcels outlined in green) is spatially anti-correlated with the LSD response map ($\rho = -0.44$, see **Methods**), suggesting that circuits modulated by LSD (i.e. serotonin, in particular 5-HT2A) may be relevant for the PC5 symptom expression. Here a serotonin receptor agonist may modulate the neuro-behavioral profile of Patient Q_{PC5} , whereas an antagonist may be effective for Patient Z_{PC5} .

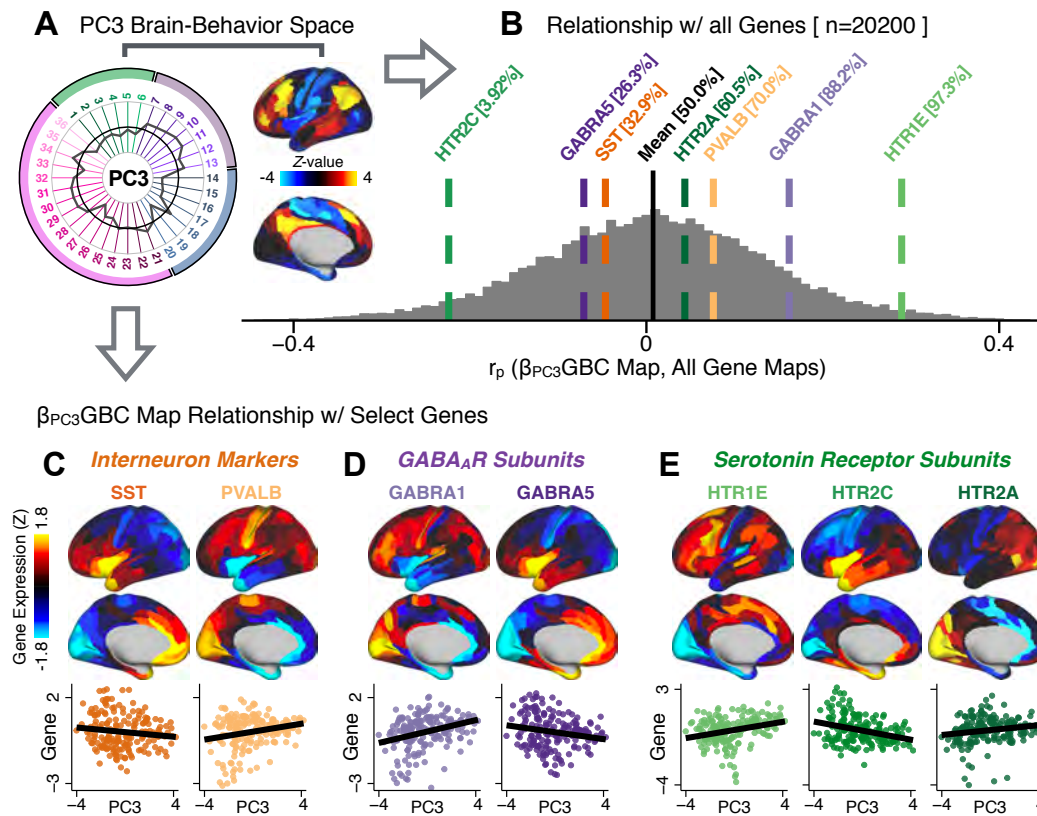


Fig. 8. Psychosis spectrum neuro-behavioral maps track neural gene expression patterns computed from the Allen Human Brain Atlas (AHBA). (A) The symptom loadings and the associated neural map jointly reflect the *PC3* brain-behavioral space (BBS) profile, which can be quantitatively related to human cortical gene expression patterns obtained from the AHBA (31). (B) Distribution of correlation values between the *PC3* BBS map and ~20,000 gene expression maps derived from the AHBA dataset. Specifically, AHBA gene expression maps were obtained using DNA microarrays from six postmortem brains, capturing gene expression topography across cortical areas. These expression patterns were then mapped onto the cortical surface models derived from the AHBA subjects' anatomical scans and aligned with the Human Connectome Project (HCP) atlas, described in prior work and methods (31). Note that because no significant inter-hemispheric differences were found in cortical gene expression all results were symmetrized to the left hemisphere, resulting in 180 parcels. We focused on a select number of psychosis-relevant genes – namely genes coding for the serotonin and GABA receptor subunits and interneuron markers. Seven genes of interest are highlighted with dashed lines. Note that the expression pattern of HTR2C (green dashed line) is at the low negative tail of the entire distribution, i.e. highly anti-correlated with *PC3* BBS map. Conversely, GABRA1 and HTR1E are on the far positive end, reflecting a highly similar gene-to-BBS spatial pattern. (C) Upper panels show gene expression patterns for two interneuron marker genes, somatostatin (SST) and parvalbumin (PVALB). Positive (yellow) regions show areas where the gene of interest is highly expressed, whereas negative (blue) regions indicate low expression values. Lower panels highlight all gene-to-BBS map spatial correlations where each value is a symmetrized cortical parcel (180 in total) from the HCP atlas parcellation. (D) Gene expression maps and spatial correlations with the *PC3* BBS map for two GABA_A receptor subunit genes: GABRA1 and GABRA5. (E) Gene expression maps and spatial correlations with the *PC3* BBS map for three serotonin receptor subunit genes: HTR1E, HTR2C, and HTR2A.

logical and gene expression neural maps that are directly relevant for PSD neurobiology.

Deriving an Individually Predictive Low-Dimensional Symptom Representation Across the Psychosis Spectrum. Psychosis spectrum is associated with notable clinical heterogeneity such deficits in cognition as well as altered beliefs (i.e. delusions), perception (i.e. hallucinations), and affect (i.e. negative symptoms) (24). This heterogeneity is captured by clinical instruments that quantify PSD symptoms across dozens of specific questions and ratings. This yields a high-dimensional symptom space that is intractable for reliable mapping of neural circuits (40). Here we show that a low-rank solution captures principal axes of PSD symptom variation, a finding in line with prior work in schizophrenia (18–24, 48).

These results highlights two key observations: i) Existing symptom reduction studies (even those in schizophrenia

specifically) have not evaluated solutions that include cognitive impairment – a hallmark deficit across the psychosis spectrum (17). Here we show that cognitive function captures a notable portion of the symptom variance independent of other axes. We observed that cognitive variation captured 10% of PSD sample variance even after accounting for ‘General’ psychopathology. ii) No study has quantified cognitive deficit variation via dimensionality reduction across multiple PSD diagnoses along with core psychosis symptoms. We found that cognitive deficits load across several PCs, but the pattern of loading was particularly evident for executive function sub-scores on certain axes (e.g. *PC5* symptom axis solution was not stable if a single composite cognitive score was used). While existing studies evaluated stability of data-reduced solutions within a single DSM category (18, 24, 49, 50), current results show that lower-dimensional PSD symptoms solutions can be reproducibly obtained across DSM diagnoses. For each data-reduced axis

some PSD patients received a score near zero. This does not imply that these patients were unimpaired; rather, the symptom configurations for these patients were orthogonal to variation along this specific symptom axis. The observation that PSD is associated with multiple complex symptom dimensions highlights an important intuition that may extend to other mental health spectra. Also, the PSD symptom axes reported here are neither definitive nor exhaustive. In fact, close to 50% of all clinical variance was not captured by the symptom PCA – an observation often overlooked in symptom data-reduction studies, which focus on attaining ‘predictive accuracy’. Such studies rarely consider how much variance remains unexplained in the final data-reduced model and, relatedly, if the proportion of explained variance is reproducible across samples. This is a key property for reliable symptom-to-mapping. Thus, we tested if this replicable low-dimensional PSD symptom space robustly mapped onto neural circuit patterns.

Leveraging a Robust Low-Dimensional Symptom Representation for Mapping Brain-Behavior Relationships. We show that the dimensionality-reduced symptom space improved the mapping onto neural circuit features (i.e. GBC), as compared to a *priori* item-level clinical symptoms (Fig. 2). This symptom-to-neural mapping was highly reproducible across various cross-validation procedures, including split-half replication (Fig. 3). The observed statistical symptom-to-neural improvement after dimensionality reduction suggests that data-driven clinical variation more robustly covaried with neural features. As noted, the low-rank symptom axes generalized across DSM diagnoses. Consequently, the mapping onto neural features (i.e. GBC) may have been restricted if only a single DSM category or clinical item was used. Importantly, as noted, traditional clinical scales are unidirectional (i.e. zero is asymptomatic, hence there is an explicit floor). Here, we show that data-driven symptom axes (e.g. PC3) were associated with bi-directional variation (i.e., no explicit floor effect). Put differently, patients who score highly on either end of these data-driven axes are severely symptomatic but in very different ways. If these axes reflect clinically meaningful phenomena at both tails, then they should more robustly map to neural feature variation, which is in line with reported effects. Therefore, by definition, the derived map for each of the PCs will reflect the neural circuitry that may be modulated by the behaviors that vary along that PC (but not others). For example, we named the PC3 axis “Psychosis Configuration” because of its strong loadings onto conventional “positive” and “negative” PSD symptoms. This PC3 “Psychosis Configuration” showed strong positive variation along neural regions that map onto the canonical default mode network (DMN), which has frequently been implicated in PSD (51–56). In turn, this bi-directional “Psychosis Configuration” axis showed strong negative variation along neural regions that map onto the sensory-motor and associative control regions, also strongly implicated in PSD (?). This ‘bi-directional’ neuro-behavioral map property may be desirable for identifying neural features that support individual patient selection.

Deriving Individually Actionable Brain-Behavior Mapping Across the Psychosis Spectrum. Deriving a neuro-behavioral mapping that is resolvable and stable at the individual patient level is a necessary benchmark for deploying symptom-to-neural ‘biomarkers’ in a clinically useful way. Therefore, there is increasing attention placed on the importance of achieving reproducibility in the psychiatric neuroimaging literature (57–60), which becomes especially important for individualized symptom-to-neural mapping. Recently, several efforts have deployed multivariate methods to quantify symptom-to-neural relationships (15, 37, 38, 61–63), highlighting how multivariate techniques may perhaps provide clinically innovative insights. However, such methods face overfitting risk for high-dimensional but underpowered datasets (25), as recently shown via formal generative modeling (40).

Here we attempted to use multivariate solutions (i.e. CCA) to quantify symptom and neural feature co-variation. In principle, CCA is theoretically well-suited to address the brain-behavioral mapping problem. However, symptom-to-neural mapping using CCA across either parcel-level or network-level solutions was not reproducible even when using a low-dimensional symptom solutions as a starting point. Therefore, while CCA (and related multivariate methods such as partial least squares) are theoretically appropriate (and regularization methods such as sparse CCA may help), in practice many available psychiatric neuroimaging datasets may not provide sufficient power to resolve stable multivariate symptom-to-neural solutions (40). A key pressing need for forthcoming studies will be to use multivariate power calculators to inform sample sizes needed for resolving stable neuro-behavioral geometries at the single subject level.

Consequently, we tested if a low-dimensional symptom solution can be used in a univariate symptom-to-neural model to optimize individually predictive features. Indeed, we found that a univariate brain-behavioral space (BBS) relationship can result in neural features that are stable for individualized prediction. Critically, we found that a patient exhibited a high PC symptom score, they were more likely to exhibit a topography of neural Δ GBC (i.e. difference relative to the group mean reference) that was topographically similar to PC symptom neural map. This suggests that optimizing such symptom-to-neural mapping solutions (and ultimately extending them to multivariate frameworks) can inform cross-diagnostic patient segmentation with respect to symptom-relevant neural features. Importantly, this could directly inform patient identification based on neural targets that are of direct symptom relevance for clinical trial design.

Utilizing Independent Molecular Neuroimaging Maps to ‘Benchmark’ Symptom-Relevant Neural Features. Selecting single patients via stable symptom-to-neural mapping BBS solutions is necessary for individual patient segmentation, which may ultimately inform treatment indication. However, are the derived symptom-to-neural maps related to a given mechanism? Here we highlight two ways to ‘benchmark’ the derived symptom-to-neural feature maps by calculating their similarity against indepen-

dent pharmacological neuroimaging and gene expression maps. We show a proof-of-principle framework for quantifying derived symptom-to-neural reference maps with two PSD-relevant neuropharmacological manipulations derived in healthy adults via LSD and ketamine. These analyses revealed that selecting single patients, via the derived symptom-to-neural mapping solution, can yield above-chance quantitative correspondence to a given molecular target map. These data highlight an important effect: it is possible to construct a “strong inference” (64) evaluation of single patients’ differential similarity to molecular target map. For instance, this approach could be applied to maps associated with already approved PSD treatments (such as clozapine, olanzapine, or chlorpromazine (65, 66)) to identify patients with symptom-to-neural configurations that best capture available treatment-covarying neural targets.

Relatedly, AHBA gene expression maps (31) may provide an a priori benchmark for treatment targets that may be associated with a given receptor profile. Here we show that identified BBS maps exhibit spatial correspondence with neural gene expression maps implicated in PSD – namely serotonin, GABA and interneuron gene expression. This gene-to-BBS mapping could be then used to select those patients that exhibit high correspondence to a given gene expression target. Collectively, this framework could inform empirically testable treatment selection methods (e.g. a patient may benefit from ketamine, but not serotonin agonists such as LSD/psilocybin). In turn, this independent molecular benchmarking framework could be extended to other approaches (e.g. positron emission tomography (PET) maps reflecting specific neural receptor density patterns (67, 68)) and iteratively optimized for quantitative patient-specific selection against actionable molecular targets.

Considerations for Generalizing Solutions Across Time, Severity and Mental Health Spectra. There are several constraints of the current result that require future optimization – namely the ability to generalize across time (early course vs. chronic patients), across a range of symptom severity (e.g. severe psychotic episode or persistent low-severity psychosis) and across distinct symptom spectra (e.g. mood). This applies to both the low-rank symptom solution and the resulting symptom-to-neural mapping. It is possible that the derived lower-dimensional symptom solution, and consequently the symptom-to-neural mapping solution, exhibits either time-dependent (i.e. state) or severity-dependent (i.e. trait) re-configuration. Relatedly, medication dose, type, and timing may also impact the solution.

Critically, these factors should constitute key extensions of an iteratively more robust model for individualized symptom-to-neural mapping across the PSD and other psychiatric spectra. Relatedly, it will be important to identify the ‘limits’ of a given BBS solution – namely a PSD-derived effect may not generalize into the mood spectrum (i.e. both the symptom space and the resulting symptom-to-neural mapping is orthogonal). It will be important to evaluate if this framework can be used to initialize symptom-to-neural mapping across other mental health symptom spectra, such as mood/anxiety

disorders.

These types questions will require longitudinal and clinically diverse studies that start prior to the onset of full-blown symptoms (e.g. the North American Prodrome Longitudinal Study (NAPLS) (69, 70)). A corollary of this point is that ~50% of unexplained symptom variance in the current PCA solution necessitates larger samples with adequate power to map this subtle but perhaps clinically essential PSD variation. Notably, the current cohort was adequately powered for symptom data reduction that drove univariate neural mapping. However, this sample was insufficiently powered for resolving stable multivariate symptom-to-neural relationships even with low-dimensional symptom features. Consequently, the limited sample size necessitated choices for dimensionality-reduction of the neural feature space in this study even for univariate analyses. While both parcellation and GBC constitute principled choices, symptom-relevant neural information may have been lost (which may be imbedded in a higher-dimensional space). One obvious solution is to increase sample sizes (e.g. via datasets such as the UK Biobank (71)). However, in parallel, it will be critical to develop neurobiologically-informed feature space reduction and/or to optimize the stability of multivariate solutions via regularization methods. Another improvement would be to optimize neural data reduction sensitivity for specific symptom variation (72).

Here we focused on the neural blood oxygen level dependent (BOLD) signal from fMRI. However, other modalities such as diffusion-weighted imaging (DWI), PET imaging, or electroencephalography (EEG) could be leveraged. Additional clinically-relevant information could be derived from genetics (such as polygenic risk scores (73–75)) or ecological momentary assessment (EMA) (76, 77), especially to parse state vs. trait biomarker variation.

Lastly, building on the proof-of-concept molecular neuroimaging comparisons, it will be imperative to eventually test such predictions in formal clinical trials. An actionable next step would be to optimize patient selection against existing treatments, which could result in higher success rates for drug development trials and potentially massive impact for developing new interventions. Critically, the opportunity to develop, validate, and refine an individual-level quantitative framework could deliver a more rapid and cost-effective way of pairing patients with the right treatments.

Conclusions. We show that complex and highly heterogeneous PSD symptom variation can be robustly reduced into a low-rank symptom solution that is cross-diagnostic, individually predictive, generalizable and incorporates cognitive deficits. In turn, the derived PSD symptom axes robustly mapped onto distinct yet replicable neural patterns, which were predictive at the single-patient level. Leveraging these stable results, we show a proof-of-principle framework for relating the derived symptom-relevant neural maps, at the individual patient level, with molecular targets implicated in PSD via LSD and ketamine neuro-pharmacological manipulations. Lastly, we used AHBA gene expression maps to show that identified PSD symptom-relevant neural maps co-

vary with serotonin, GABA and interneuron neural gene expression patterns. This set of symptom-to-neural mapping results can be iteratively and quantitatively optimized for personalized treatment segmentation endpoints.

Competing Interests. J.L.J., C.F., J.B.B., G.R., Z.T., J.D., and J.D.M. consult for BlackThorn Therapeutics. A.A. consults and is a SAB member for BlackThorn Therapeutics Inc. W.J.M. is an employee of The Janssen Pharmaceutical Companies of Johnson & Johnson. J.L.J., J.D.M., and A.A. are co-inventors for the following pending patent: Anticevic A, Murray JD, Ji JL: Systems and Methods for Neuro-Behavioral Relationships in Dimensional Geometric Embedding (N-BRIDGE), PCT International Application No. PCT/US2119/022110, filed March 13, 2019. J.D.M., A.A., and W.J.M. are co-inventors for the following pending patent: Murray JD, Anticevic A, Martin, WJ: Methods and tools for detecting, diagnosing, predicting, prognosticating, or treating a neurobehavioral phenotype in a subject, U.S. Application No. 16/149,903 filed on October 2, 2018, U.S. Application or PCT International Application No. 18/054,009 filed on October 2, 2018.

ACKNOWLEDGEMENTS

We thank J.H. Krystal and I. Levy for useful discussions. This work was supported by the Heffter Research Institute (Grant No. 1-190420), Swiss Neuromatrix Foundation (Grant No. 2016-0111), Swiss National Science Foundation under the framework of Neuron Cofund (Grant No. 01EW1908), Slovenian Research Agency (ARRS) (Grant Nos. J7-8275, J7-6829, P3-0338 to GR), 1U01MH121766 (to JHK & AA), R01MH112746 (to JDM), SFARI Pilot Award (to JDM & AA), 5R01MH12189 (to AA), 5R01MH108590 (to AA), and by BlackThorn Therapeutics (J.D.M. & A.A.). Data were provided by the Bipolar-Schizophrenia Network on Intermediate Phenotypes (BSNIP) consortium via the National Institute of Mental Health Data Archive (NDA).

Bibliography

- Mauricio Tohen, Eduard Vieta, Joseph Calabrese, Terence A Ketter, Gary Sachs, Charles Bowden, Philip B Mitchell, Franca Centorrino, Richard Risper, Robert W Baker, et al. Efficacy of olanzapine and olanzapine-fluoxetine combination in the treatment of bipolar I depression. *Archives of general psychiatry*, 60(11):1079–1088, 2003.
- Joseph P McEvoy, Jeffrey A Lieberman, Diana O Perkins, Robert M Hamer, Hongbin Gu, Arthur Lazarus, Dennis Switzer, Christina Olexy, Peter Weiden, and Stephen D Strakowski. Efficacy and tolerability of olanzapine, quetiapine, and risperidone in the treatment of early psychosis: a randomized, double-blind 52-week comparison. *American journal of Psychiatry*, 164(7):1050–1060, 2007.
- Marije aan het Rot, Katherine A Collins, James W Murrough, Andrew M Perez, David L Reich, Dennis S Charney, and Sanjay J Mathew. Safety and efficacy of repeated-dose intravenous ketamine for treatment-resistant depression. *Biological psychiatry*, 67(2):139–145, 2010.
- Andrea Cipriani, Toshi A Furukawa, Georgia Salanti, Anna Chaimani, Lauren Z Atkinson, Yusuke Ogawa, Stefan Leucht, Henricus G Ruhe, Erick H Turner, Julian PT Higgins, et al. Comparative efficacy and acceptability of 21 antidepressant drugs for the acute treatment of adults with major depressive disorder: a systematic review and network meta-analysis. *Focus*, 16(4):420–429, 2018.
- American Psychiatric Association and American Psychiatric Association. Task Force on DSM-IV. *Diagnostic and Statistical Manual of Mental Disorders: DSM-IV. DIAGNOSTIC AND STATISTICAL MANUAL OF MENTAL DISORDERS*. American Psychiatric Association, 1994. ISBN 9780890420621.
- Katharine A Phillips, Michael B First, and Harold Alan Pincus. *Advancing DSM: Dilemmas in psychiatric diagnosis*. American Psychiatric Pub, 2008.
- Seth J Gillihan and Erik Parens. Should we expect "neural signatures" for dsm diagnoses? *Journal of Clinical Psychiatry*, 72(10):1383, 2011.
- Alan Anticevic, Michael W Cole, Grega Repovs, John D Murray, Margaret S Brumbaugh, Anderson M Winkler, Aleksandar Savic, John H Krystal, Godfrey D Pearson, and David C Glahn. Characterizing thalamo-cortical disturbances in schizophrenia and bipolar illness. *Cerebral cortex*, 24(12):3116–3130, 2013.
- BJ Casey, Nick Craddock, Bruce N Cuthbert, Steven E Hyman, Francis S Lee, and Kerry J Ressler. Dsm-5 and rdco: progress in psychiatry research? *Nature Reviews Neuroscience*, 14(11):810, 2013.
- Thomas R Insel. The nimh research domain criteria (rdco) project: precision medicine for psychiatry. *American Journal of Psychiatry*, 171(4):395–397, 2014.
- Carol A Tamminga, Godfrey Pearson, Matcheri Keshavan, John Sweeney, Brett Clementz, and Guntav Thaker. Bipolar and schizophrenia network for intermediate phenotypes: outcomes across the psychosis continuum. *Schizophrenia bulletin*, 40(Suppl_2):S131–S137, 2014.
- Nikolaos Koutsouleris, Lana Kambeitz-Illankovic, Stephan Ruhrmann, Marlene Rosen, Anne Ruef, Dominic B Dwyer, Marco Paolini, Katharine Chisholm, Joseph Kambeitz, Theresa Haidl, et al. Prediction models of functional outcomes for individuals in the clinical high-risk state for psychosis or with recent-onset depression: a multimodal, multisite machine learning analysis. *JAMA psychiatry*, 75(11):1156–1172, 2018.
- BJ Casey, Tariq Cannonier, May I Conley, Alexandra O Cohen, Deanna M Barch, Mary M Heitzeg, Mary E Soules, Theresa Teslovich, Danielle V Dellarco, Hugh Garavan, et al. The adolescent brain cognitive development (abdc) study: imaging acquisition across 21 sites. *Developmental cognitive neuroscience*, 32:43–54, 2018.
- Adriana Di Martino, Chao-Gan Yan, Qingyang Li, Erin Denio, Francisco X Castellanos, Kaat Alaerts, Jeffrey S Anderson, Michal Assaf, Susan Y Bookheimer, Mirella Dapretto, et al. The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular psychiatry*, 19(6):659, 2014.
- Andrew T Drysdale, Logan Grosenick, Jonathan Downar, Katharine Dunlop, Farrokh Mansouri, Yue Meng, Robert N Fetcho, Benjamin Zebly, Desmond J Oathes, Amit Etkin, et al. Resting-state connectivity biomarkers define neurophysiological subtypes of depression. *Nature medicine*, 23(1):28, 2017.
- Lee J Cronbach and Paul E Mehl. Construct validity in psychological tests. *Psychological bulletin*, 52(4):281, 1955.
- Deanna M Barch, Juan Bustillo, Wolfgang Gaebel, Raquel Gur, Stephan Heckers, Dolores Malaspina, Michael J Owen, Susan Schultz, Rajiv Tandon, Ming Tsuang, et al. Logic and justification for dimensional assessment of symptoms and related clinical phenomena in psychosis: relevance to dsm-5. *Schizophrenia research*, 150(1):15–20, 2013.
- Mark van der Gaag, Anke Cuijpers, Tonko Hoffman, Mila Remijsen, Ron Hijman, Lieuwe de Haan, Berno van Meijel, Peter N van Harten, Lucia Valmaggia, Marc De Hert, et al. The five-factor model of the positive and negative syndrome scale i: confirmatory factor analysis fails to confirm 25 published five-factor solutions. *Schizophrenia research*, 85(1-3):273–279, 2006.
- Jean-Pierre Lindenmayer, Ruth Bernstein-Hyman, and Sandra Grochowski. A new five factor model of schizophrenia. *Psychiatric Quarterly*, 65(4):299–322, 1994.
- Robin Emsley, Jonathan Rabinowitz, Martijn Torremans, RIS-INT-35 Early Psychosis Global Working Group, et al. The factor structure for the positive and negative syndrome scale (panss) in recent-onset psychosis. *Schizophrenia research*, 61(1):47–57, 2003.
- Sonia Dollfus, Brian Everitt, Jean Marie Ribeyre, Francoise Assouly-Besse, Charlie Sharp, and Michel Petit. Identifying subtypes of schizophrenia by cluster analyses. *Schizophrenia bulletin*, 22(3):545–555, 1996.
- Jack J Blanchard and Alex S Cohen. The structure of negative symptoms within schizophrenia: implications for assessment. *Schizophrenia bulletin*, 32(2):238–245, 2006.
- Ji Chen, Kaustubh R Patil, Susanne Weis, Kang Sim, Thomas Nickl-Jockschat, Juan Zhou, André Aleman, Iris E Sommer, Edith J Liemburg, Felix Hoffstaedter, et al. Neurobiological divergence of the positive and negative schizophrenia subtypes identified on a new factor structure of psychopathology using non-negative factorization: An international machine learning study. *Biological psychiatry*, 2019.
- Jérémy Lefort-Besnard, Gaël Varoquaux, Birgit Derntl, Oliver Gruber, Andre Aleman, Renaud Jardri, Iris Sommer, Bertrand Thirion, and Danilo Bzdok. Patterns of schizophrenia symptoms: hidden structure in the panss questionnaire. *Transl Psychiatry*, 8(1):237, 10 2018. doi: 10.1038/s41398-018-0294-4.
- Richard Dinga, Lianne Schmaal, Brenda WJH Penninx, Marie Jose van Tol, Dick J Veltman, Laura van Velzen, Maarten Mennes, Nic JA van der Wee, and Andre F Marquand. Evaluating the evidence for biotypes of depression: Methodological replication and extension of. *NeuroImage: Clinical*, 22:101796, 2019.
- Carla M Canuso, Cynthia A Bossie, Young Zhu, Eriene Youssef, and David L Dunner. Psychotic symptoms in patients with bipolar mania. *Journal of Affective Disorders*, 111(2-3):164–169, 2008.
- Stanley R Kay, Abraham Fiszbein, and Lewis A Opler. The positive and negative syndrome scale (panss) for schizophrenia. *Schizophrenia bulletin*, 13(2):261–276, 1987.
- Katrin H Preller, Joshua B Burt, Jie Lisa Ji, Charles H Schleifer, Brendan D Adkinson, Philipp Stämpfli, Erich Seifritz, Grega Repovs, John H Krystal, John D Murray, et al. Changes in global and thalamic brain connectivity in lsd-induced altered states of consciousness are attributable to the 5-HT_{2A} receptor. *Elife*, 7:e35082, 2018.
- Alan Anticevic, Philip R Corlett, Michael W Cole, Aleksandar Savic, Mark Gancsos, Yanqing Tang, Grega Repovs, John D Murray, Naomi R Driesen, Peter T Morgan, et al. N-methyl-D-aspartate receptor antagonist effects on prefrontal cortical connectivity better model early than chronic schizophrenia. *Biological psychiatry*, 77(6):569–580, 2015.
- Michael J Hawrylycz, Ed S Lein, Angela L Guillozet-Bongaarts, Elaine H Shen, Lydia Ng, Jeremy A Miller, Louie N Van De Lagemaat, Kimberly A Smith, Amanda Ebbert, Zackery L Riley, et al. An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature*, 489(7416):391, 2012.
- Joshua B Burt, Murat Demirtas, William J Eckner, Natasha M Navejar, Jie Lisa Ji, William J Martin, Alberto Bernacchia, Alan Anticevic, and John D Murray. Hierarchy of transcriptomic specialization across human cortex captured by structural neuroimaging topography. *Nature neuroscience*, 21(9):1251, 2018.
- Matcheri S Keshavan, David W Morris, John A Sweeney, Godfrey Pearson, Guntav Thaker, Larry J Seidman, Shaun M Eack, and Carol Tamminga. A dimensional approach to the psychosis spectrum between bipolar disorder and schizophrenia: the schizo-bipolar scale. *Schizophrenia research*, 133(1-3):250–254, 2011.
- Jie Lisa Ji, Marjolein Spronk, Kaustubh Kulkarni, Grega Repovs, Alan Anticevic, and Michael W Cole. Mapping the human brain's cortical-subcortical functional network organization. *NeuroImage*, 185:35–57, 2019.
- Michael W Cole, Sudhir Pathak, and Walter Schneider. Identifying the brain's most globally connected regions. *NeuroImage*, 49(4):3132–3148, 2010.
- Michael D Fox, Abraham Z Snyder, Justin L Vincent, Maurizio Corbetta, David C Van Essen, and Marcus E Raichle. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, 102

- (27):9673–9678, 2005.
36. Matthew F Glasser, Timothy S Coalson, Emma C Robinson, Carl D Hacker, John Harwell, Essa Yacoub, Kamil Ugurbil, Jesper Andersson, Christian F Beckmann, Mark Jenkinson, et al. A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615):171, 2016.
 37. Cedric Huchuan Xia, Zongming Ma, Rastko Ciric, Shi Gu, Richard F Betzel, Antonia N Kaczurkin, Monica E Calkins, Philip A Cook, Angel Garcia de la Garza, Simon N Vandekar, et al. Linked dimensions of psychopathology and connectivity in functional brain networks. *Nature communications*, 9(1):3003, 2018.
 38. Meichen Yu, Kristin A Linn, Russell T Shinohara, Desmond J Oathes, Philip A Cook, Romain Duprat, Tyler M Moore, Maria A Oquendo, Mary L Phillips, Melvin McInnis, et al. Childhood trauma history is linked to abnormal brain connectivity in major depression. *Proceedings of the National Academy of Sciences*, 116(17):8582–8590, 2019.
 39. David R Hardoon, Sandor Szedmak, and John Shawe-Taylor. Canonical correlation analysis: An overview with application to learning methods. *Neural computation*, 16(12):2639–2664, 2004.
 40. Markus Helmer, Shaun D Warrington, Ali-Reza Mohammadi-Nejad, Jie Lisa Ji, Amber Howell, Benjamin Rosand, Alan Anticevic, Stamatis N Sotiropoulos, and John D Murray. On stability of canonical correlation analysis and partial least squares with application to brain-behavior associations. *bioRxiv*, 2020. doi: 10.1101/2020.08.25.265546.
 41. Robert M Berman, Angela Capiello, Amit Anand, Dan A Oren, George R Heninger, Dennis S Charney, and John H Krystal. Antidepressant effects of ketamine in depressed patients. *Biological psychiatry*, 47(4):351–354, 2000.
 42. John H Krystal, Laurence P Karper, John P Seibyl, Glenna K Freeman, Richard Delaney, J Douglas Bremner, George R Heninger, Malcolm B Bowers, and Dennis S Charney. Subanesthetic effects of the noncompetitive nmda antagonist, ketamine, in humans: psychotomimetic, perceptual, cognitive, and neuroendocrine responses. *Archives of general psychiatry*, 51(3):199–214, 1994.
 43. Javier González-Maeso, Noelia V Weisstaub, Mingming Zhou, Pokman Chan, Lidija Ivic, Rosalind Ang, Alena Lira, Maria Bradley-Moore, Yongchao Ge, Qiang Zhou, et al. Hal-lucinogens recruit specific cortical 5-HT_{2A} receptor-mediated signaling pathways to affect behavior. *Neuron*, 53(3):439–452, 2007.
 44. Christina T Egan, Katharine Herrick-Davis, Keith Miller, Richard A Glennon, and Milt Teitler. Agonist activity of lsd and lisuride at cloned 5HT_{2A} and 5HT_{2C} receptors. *Psychopharmacology*, 136(4):409–414, 1998.
 45. Bitu Moghaddam and Daniel Javitt. From revolution to evolution: the glutamate hypothesis of schizophrenia and its implication for treatment. *Neuropsychopharmacology*, 37(1):4–15, 2012.
 46. Alan Anticevic, Mark Gancsos, John D. Murray, Grega Repovs, Naomi R. Driesen, Debra J. Ennis, Mark J. Niciu, Peter T. Morgan, Toral S. Surti, Michael H. Bloch, Ramachandran Ramani, Mark A. Smith, Xiao-Jing Wang, John H. Krystal, and Philip R. Corlett. Nmda receptor function in large-scale anticorrelated neural systems with implications for cognition and schizophrenia. *Proceedings of the National Academy of Sciences*, 109(41):16720–16725, 2012. ISSN 0027-8424. doi: 10.1073/pnas.1208494109.
 47. Sean L Kitson. 5-hydroxytryptamine (5-HT) receptor ligands. *Current pharmaceutical design*, 13(25):2621–2637, 2007.
 48. Leonard White, Philip D Harvey, Lewis Opler, and JP Lindenmayer. Empirical assessment of the factorial structure of clinical symptoms in schizophrenia. *Psychopathology*, 30(5):263–274, 1997.
 49. Mark van der Gaag, Tonko Hoffman, Mila Remijsen, Ron Hijman, Lieuwe de Haan, Berno van Meijel, Peter N van Harten, Lucia Valmaggia, Marc De Hert, Anke Cuijpers, et al. The five-factor model of the positive and negative syndrome scale ii: a ten-fold cross-validation of a revised model. *Schizophrenia research*, 85(1-3):280–287, 2006.
 50. Jean-Pierre Lindenmayer, Sandra Grochowski, and Ruth B Hyman. Five factor model of schizophrenia: replication across samples. *Schizophrenia research*, 14(3):229–234, 1995.
 51. Susanna L Fryer, Scott W Woods, Kent A Kiehl, Vince D Calhoun, Godfrey Pearlson, Brian J Roach, Judith M Ford, Vinod H Srihari, Thomas H McGlashan, and Daniel H Mathalon. Deficient suppression of default mode regions during working memory in individuals with early psychosis and at clinical high-risk for psychosis. *Frontiers in psychiatry*, 4:92, 2013.
 52. Alan Anticevic, Michael W Cole, John D Murray, Philip R Corlett, Xiao-Jing Wang, and John H Krystal. The role of default network deactivation in cognition and disease. *Trends in cognitive sciences*, 16(12):584–592, 2012.
 53. Dost Öngür, Miriam Lundy, Ian Greenhouse, Ann K Shinn, Vinod Menon, Bruce M Cohen, and Perry F Renshaw. Default mode network abnormalities in bipolar disorder and schizophrenia. *Psychiatry Research: Neuroimaging*, 183(1):59–68, 2010.
 54. Neil D Woodward, Baxter Rogers, and Stephan Heckers. Functional resting-state networks are differentially affected in schizophrenia. *Schizophrenia research*, 130(1-3):86–93, 2011.
 55. Justin T Baker, Avram J Holmes, Grace A Masters, BT Thomas Yeo, Fenna Krienen, Randy L Buckner, and Dost Öngür. Disruption of cortical association networks in schizophrenia and psychotic bipolar disorder. *JAMA psychiatry*, 71(2):109–118, 2014.
 56. Shashwath A Meda, Gualberto Ruaño, Andreas Windemuth, Kasey O’Neil, Clifton Berwise, Sabra M Dunn, Leah E Boccaccio, Balaji Narayanan, Mohan Kocherla, Emma Sprooten, et al. Multivariate analysis reveals genetic associations of the resting default mode network in psychotic bipolar disorder and schizophrenia. *Proceedings of the National Academy of Sciences*, 111(19):E2066–E2075, 2014.
 57. Stephanie Noble, Dustin Scheinost, and R Todd Constable. A decade of test-retest reliability of functional connectivity: A systematic review and meta-analysis. *Neuroimage*, 203:116157, 2019.
 58. Joshua H Balsters, Dante Mantini, Matthew AJ Apps, Simon B Eickhoff, and Nicole Wenderoth. Connectivity-based parcellation increases network detection sensitivity in resting state fmri: an investigation into the cingulate cortex in autism. *NeuroImage: Clinical*, 11:494–507, 2016.
 59. Hengyi Cao, Sarah C McEwen, Jennifer K Forsyth, Dylan G Gee, Carrie E Bearden, Jean Addington, Bradley Goodyear, Kristin S Cadenhead, Helene Mirzakhani, Barbara A Cornblatt, et al. Toward leveraging human connectomic data in large consortia: generalizability of fmri-based brain graphs across sites, sessions, and paradigms. *Cerebral Cortex*, 29(3):1263–1279, 2019.
 60. Choong-Wan Woo, Luke J Chang, Martin A Lindquist, and Tor D Wager. Building better biomarkers: brain models in translational neuroimaging. *Nature neuroscience*, 20(3):365, 2017.
 61. Stephen M Smith, Thomas E Nichols, Diego Vidaurre, Anderson M Winkler, Timothy EJ Behrens, Matthew F Glasser, Kamil Ugurbil, Deanna M Barch, David C Van Essen, and Karla L Miller. A positive-negative mode of population covariation links brain connectivity, demographics and behavior. *Nature neuroscience*, 18(11):1565, 2015.
 62. Dominik A Moser, Gaelle E Doucet, Won Hee Lee, Alexander Rasgon, Hannah Krinsky, Evan Leib, Alex Ing, Gunter Schumann, Natalie Rasgon, and Sophia Frangou. Multivariate associations among behavioral, clinical, and multimodal imaging phenotypes in patients with psychosis. *JAMA psychiatry*, 75(4):386–395, 2018.
 63. Amanda L Rodrigue, Jennifer E McDowell, Neeraj Tandon, Matcheri S Keshavan, Carol A Tamminga, Godfrey D Pearlson, John A Sweeney, Robert D Gibbons, and Brett A Clementz. Multivariate relationships between cognition and brain anatomy across the psychosis spectrum. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 3(12):992–1002, 2018.
 64. John R Platt. Strong inference. *science*, 146(3642):347–353, 1964.
 65. John Lally and James H MacCabe. Antipsychotic medication in schizophrenia: a review. *British medical bulletin*, 114(1):169–179, 2015.
 66. S Miyamoto, GE Duncan, CE Marx, and JA Lieberman. Treatments for schizophrenia: a critical review of pharmacology and mechanisms of action of antipsychotic drugs. *Molecular psychiatry*, 10(1):79–104, 2005.
 67. Lars Farde, Håkan Hall, Erling Ehrin, and Göran Sedvall. Quantitative analysis of d2 dopamine receptor binding in the living human brain by pet. *Science*, 231(4735):258–261, 1986.
 68. Ryosuke Arakawa, Hiroshi Ito, Akihiro Takano, Hidehiko Takahashi, Takuya Morimoto, Takeshi Sassa, Katsuya Ohta, Motoichiro Kato, Yoshiro Okubo, and Tetsuya Suhara. Dose-finding study of paliperidone er based on striatal and extrastriatal dopamine d2 receptor occupancy in patients with schizophrenia. *Psychopharmacology*, 197(2):229–235, 2008.
 69. Jean Addington, Kristin S Cadenhead, Barbara A Cornblatt, Daniel H Mathalon, Thomas H McGlashan, Diana O Perkins, Larry J Seidman, Ming T Tsuang, Elaine F Walker, Scott W Woods, et al. North american prodrome longitudinal study (napls 2): overview and recruitment. *Schizophrenia research*, 142(1-3):77–82, 2012.
 70. Larry J Seidman, Anthony J Giuliano, Eric C Meyer, Jean Addington, Kristin S Cadenhead, Tyrone D Cannon, Thomas H McGlashan, Diana O Perkins, Ming T Tsuang, Elaine F Walker, et al. Neuropsychology of the prodrome to psychosis in the napls consortium: relationship to family history and conversion to psychosis. *Archives of general psychiatry*, 67(6):578–588, 2010.
 71. Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, et al. Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine*, 12(3):e1001779, 2015.
 72. Zarrar Shehzad, Clare Kelly, Philip T Reiss, R Cameron Craddock, John W Emerson, Katie McMahon, David A Copland, F Xavier Castellanos, and Michael P Milham. A multivariate distance-based analytic framework for connectome-wide association studies. *Neuroimage*, 93 Pt 1:74–94, Jun 2014. doi: 10.1016/j.neuroimage.2014.02.024.
 73. Steven E Hyman. Diagnosis of mental disorders in light of modern genetics. *The conceptual evolution of DSM-5*, 3:18, 2011.
 74. Michael C O’Donovan, Nick J Craddock, and Michael J Owen. Genetics of psychosis: insights from views across the genome. *Human genetics*, 126(1):3–12, 2009.
 75. Hengyi Cao, Hang Zhou, and Tyrone D Cannon. Functional connectome-wide associations of schizophrenia polygenic risk. *Molecular Psychiatry*, pages 1–9, 2020.
 76. Cheryl M Corcoran, Facundo Carrillo, Diego Fernández-Slezak, Gillinder Bedi, Casimir Klim, Daniel C Javitt, Carrie E Bearden, and Guillermo A Cecchi. Prediction of psychosis across protocols and risk cohorts using automated language analysis. *World Psychiatry*, 17(1):67–75, 2018.
 77. Tara A Niendam, Laura M Tully, Ana-Maria Iosif, Divya Kumar, Kathleen E Nye, Jennifer C Denton, Lauren N Zaksorn, Taylor L Fedechko, and Katherine M Pierce. Enhancing early psychosis treatment using smartphone technology: a longitudinal feasibility and validity study. *Journal of Psychiatric Research*, 96:239–246, 2018.
 78. Shashwath A Meda, Zheng Wang, Elena I Ilveta, Gaurav Poudyal, Matcheri S Keshavan, Carol A Tamminga, John A Sweeney, Brett A Clementz, David J Schretlen, Vincent D Calhoun, et al. Frequency-specific neural signatures of spontaneous low-frequency resting state fluctuations in psychosis: evidence from bipolar-schizophrenia network on intermediate phenotypes (b-snip) consortium. *Schizophrenia bulletin*, 41(6):1336–1348, 2015.
 79. Julia M Sheffield, Sridhar Kandala, Carol A Tamminga, Godfrey D Pearlson, Matcheri S Keshavan, John A Sweeney, Brett A Clementz, Dov B Lerman-Sinkoff, S Kristian Hill, and Deanna M Barch. Transdiagnostic associations between functional brain network integrity and cognition. *JAMA psychiatry*, 74(6):605–613, 2017.
 80. Shashwath A Meda, Adrienne Gill, Michael C Stevens, Raymond P Lorenzoni, David C Glahn, Vince D Calhoun, John A Sweeney, Carol A Tamminga, Matcheri S Keshavan, Guntav Thaker, et al. Differences in resting-state functional magnetic resonance imaging functional network connectivity between schizophrenia and psychotic bipolar probands and their unaffected first-degree relatives. *Biological psychiatry*, 71(10):881–889, 2012.
 81. Matthew F Glasser, Stamatis N Sotiropoulos, J Anthony Wilson, Timothy S Coalson, Bruce Fischl, Jesper L Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R Polimeni, et al. The minimal preprocessing pipelines for the human connectome project. *Neuroimage*, 80:105–124, 2013.
 82. Jie Lisa Ji, Caroline Diehl, Charles Schleifer, Carol A Tamminga, Matcheri S Keshavan, John A Sweeney, Brett A Clementz, S Kristian Hill, Godfrey Pearlson, Genevieve Yang, Gina Creatura, John H Krystal, Grega Repovs, John Murray, Anderson Winkler, and Alan Anticevic. Schizophrenia Exhibits Bi-directional Brain-Wide Alterations in Cortico-Striato-Cerebellar Circuits. *Cerebral Cortex*, 06 2019. ISSN 1047-3211. doi: 10.1093/cercor/bhy306.
 83. Mark Jenkinson, Peter Bannister, Michael Brady, and Stephen Smith. Improved optimiza-

- tion for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*, 17(2):825–841, 2002.
84. Martin Reuter, Nicholas J Schmansky, H Diana Rosas, and Bruce Fischl. Within-subject template estimation for unbiased longitudinal image analysis. *Neuroimage*, 61(4):1402–1418, 2012.
 85. Jonathan D Power, Kelly Anne Barnes, Abraham Z Snyder, Bradley L Schlaggar, and Steven E Petersen. Steps toward optimizing motion artifact removal in functional connectivity mri; a reply to carp. *Neuroimage*, 76, 2013.
 86. Alan Anticevic, Mark Gancsos, John D Murray, Grega Repovs, Naomi R Driesen, Debra J Ennis, Mark J Niciu, Peter T Morgan, Toral S Surti, Michael H Bloch, et al. Nmda receptor function in large-scale anticorrelated neural systems with implications for cognition and schizophrenia. *Proceedings of the National Academy of Sciences*, 109(41):16720–16725, 2012.
 87. Jonathan D Power, Mark Plitt, Stephen J Gotts, Prantik Kundu, Valerie Voon, Peter A Bandettini, and Alex Martin. Ridding fmri data of motion-related influences: Removal of signals with distinct spatial and physical bases in multiecho data. *Proceedings of the National Academy of Sciences*, 115(9):E2105–E2114, 2018.
 88. Jonathan D Power, Mark Plitt, Timothy O Laumann, and Alex Martin. Sources and implications of whole-brain fmri signals in humans. *Neuroimage*, 146:609–625, 2017.
 89. Genevieve J Yang, John D Murray, Matthew Glasser, Godfrey D Pearson, John H Krystal, Charlie Schleifer, Grega Repovs, and Alan Anticevic. Altered global signal topography in schizophrenia. *Cerebral cortex*, 27(11):5156–5169, 2016.
 90. Matthew F Glasser, Timothy S Coalson, Janine D Bijsterbosch, Samuel J Harrison, Michael P Harms, Alan Anticevic, David C Van Essen, and Stephen M Smith. Using temporal ica to selectively remove global noise while preserving global signal in functional mri data. *NeuroImage*, 181:692–717, 2018.
 91. Kevin M Aquino, Ben D Fulcher, Linden Parkes, Kristina Sabarwal, and Alex Fornito. Identifying and removing widespread signal deflections from fmri data: Rethinking the global signal regression problem. *bioRxiv*, page 662726, 2019.
 92. Richard SE Keefe, Terry E Goldberg, Philip D Harvey, James M Gold, Margaret P Poe, and Leigh Coughenour. The brief assessment of cognition in schizophrenia: reliability, sensitivity, and comparison with a standard neurocognitive battery. *Schizophrenia research*, 68(2-3):283–297, 2004.
 93. Carol A. Tamminga, Elena I. Ivleva, Matcheri S. Keshavan, Godfrey D. Pearson, Brett A. Clementz, Bradley Witte, David W. Morris, Jeffrey Bishop, Gunvant K. Thaker, and John A. Sweeney. Clinical phenotypes of psychosis in the bipolar-schizophrenia network on intermediate phenotypes (b-snip). *American Journal of Psychiatry*, 170(11):1263–1274, 2013. doi: 10.1176/appi.ajp.2013.12101339. PMID: 23846857.
 94. Michael W Cole, Genevieve J Yang, John D Murray, Grega Repovs, and Alan Anticevic. Functional connectivity change as shared signal dynamics. *Journal of neuroscience methods*, 259:22–39, 2016.
 95. Stephen M Smith and Thomas E Nichols. Threshold-free cluster enhancement: addressing problems of smoothing, threshold dependence and localisation in cluster inference. *Neuroimage*, 44(1):83–98, 2009.
 96. Anderson M Winkler, Gerard R Ridgway, Matthew A Webster, Stephen M Smith, and Thomas E Nichols. Permutation inference for the general linear model. *Neuroimage*, 92:381–397, 2014.
 97. Tyler M Moore, Steven P Reise, Raquel E Gur, Hakon Hakonarson, and Ruben C Gur. Psychometric properties of the penn computerized neurocognitive battery. *Neuropsychology*, 29(2):235, 2015.
 98. Ruben C Gur, Jan Richard, Paul Hughett, Monica E Calkins, Larry Macy, Warren B Bilker, Colleen Brensinger, and Raquel E Gur. A cognitive neuroscience-based computerized battery for efficient measurement of individual differences: standardization and initial construct validation. *Journal of neuroscience methods*, 187(2):254–262, 2010.
 99. David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, Wu-Minn HCP Consortium, et al. The wu-minn human connectome project: an overview. *Neuroimage*, 80:62–79, 2013.
 100. A. Hyvarinen. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks*, 10(3):626–634, May 1999. ISSN 1045-9227. doi: 10.1109/72.761722.
 101. Nathaniel E. Helwig and Sungjin Hong. A critique of tensor probabilistic independent component analysis: Implications and recommendations for multi-subject fmri data analysis. *Journal of Neuroscience Methods*, 213(2):263–273, 3 2013. ISSN 0165-0270. doi: 10.1016/j.jneumeth.2012.12.009.
 102. Alan J Gelenberg. The catatonic syndrome. *The Lancet*, 307(7973):1339–1341, 1976.
 103. Rebekka Lencer, Andreas Sprenger, James L Reilly, Jennifer E McDowell, Leah H Rubin, Judith A Badner, Matcheri S Keshavan, Godfrey D Pearson, Carol A Tamminga, Elliot S Gershon, et al. Pursuit eye movements as an intermediate phenotype across psychotic disorders: evidence from the b-snip study. *Schizophrenia research*, 169(1-3):326–333, 2015.

Methods

Overall Data Collection and Study Design. We used publicly available behavioral and neural data from the Bipolar-Schizophrenia Network on Intermediate Phenotypes (BSNIP) consortium (11). All data were obtained from the National Data Archive (NDA) repository (https://nda.nih.gov/edit_collection.html?id=2274). Participants were collected at 6 sites across the United States.

Full recruitment details are provided in prior publications (11, 78, 79). In brief, participants were recruited through advertising in Baltimore MD, Boston MA, Chicago IL, Dallas TX, and Hartford CT. All assessments were standardized across sites. Participants were excluded if they had i) a history of seizures or head injury resulting in >10 minutes loss of consciousness, ii) positive drug screen on the day of testing, iii) a diagnosis of substance abuse in the past 30 days or substance dependence in the past 6 months, iv) history of serious medical or neurological disorder that would likely affect cognitive functioning, v), history of serious medical or neurological disorder that would likely affect cognitive functioning, vi) insufficient English proficiency, or vii) an age-corrected Wide-Range Achievement Test (4th edition) reading test standard score <65. Additionally, participants were required to have had no change in medication and been clinically stable over the past month. Participants completed a SCID interview and were given a diagnosis via consensus from study clinicians; participants with an Axis I clinical psychosis diagnosis were additionally given assessments including the Positive and Negative Symptom Scale (PANSS)(27). Note that apart from the measures used in this paper, the full BSNIP dataset includes a rich characterization of measures from multiple modalities, including electrophysiological, eye tracking, structural and diffusion neuroimaging, as well as a number of clinical batteries, which are not quantified in this study. We used data from a total of 638 participants with complete behavioral and neural data after preprocessing and quality control, including 202 healthy controls (CON), 167 patients diagnosed with schizophrenia (SZP), 119 patients with schizoaffective disorder (SADP), and 150 patients with bipolar disorder with psychosis (BPP). For full demographic and clinical characteristics of the sample see **Table S1**.

Neural Data Acquisition and Preprocessing. Participants completed a neural magnetic resonance imaging (MRI) scan at 3T, including resting-state functional blood-oxygen-level-dependent imaging (BOLD) and a magnetization-prepared rapid gradient-echo (MP-RAGE) sequence for T1 weighted data. Full details on scanners and acquisition parameters used at each of the sites have previously been described (80) and are summarized here in **Table S2**. Neuroimaging data were preprocessed using the Human Connectome Project (HCP) minimal preprocessing pipeline (81), adapted for compatibility with “legacy” data, which are now featured as a standard option in the HCP pipelines provided by our team (<https://github.com/Washington-University/HCPpipelines/pull/156>). These modifications to the HCP pipelines were necessary as the BSNIP data did not include a standard field map and did not incorporate a T2w high-resolution image. The adaptations for single-band BOLD acquisition have previously been described in detail (82).

A summary of the HCP pipelines is as follows: the T1-weighted structural images were first aligned by warping them to the standard Montreal Neurological Institute-152 (MNI-152) brain template in a single step, through a combi-

nation of linear and non-linear transformations via the FMRIB Software Library (FSL) linear image registration tool (FLIRT) and non-linear image registration tool (FNIRT) (83). Next, FreeSurfer's recon-all pipeline was used to segment brain-wide gray and white matter to produce individual cortical and subcortical anatomical segmentations (84). Cortical surface models were generated for pial and white matter boundaries as well as segmentation masks for each subcortical grey matter voxel. Using the pial and white matter surface boundaries, a 'cortical ribbon' was defined along with corresponding subcortical voxels, which were combined to generate the neural file in the Connectivity Informatics Technology Initiative (CIFTI) volume/surface 'grayordinate' space for each individual subject (81). BOLD data were motion-corrected by aligning to the middle frame of every run via FLIRT in the initial NIFTI volume space. In turn, a brain-mask was applied to exclude signal from non-brain tissue. Next, cortical BOLD data were converted to the CIFTI gray matter matrix by sampling from the anatomically-defined gray matter cortical ribbon and subsequently aligned to the HCP atlas using surface-based nonlinear deformation (81). Subcortical voxels were aligned to the MNI-152 atlas using whole-brain non-linear registration and then the FreeSurfer-defined subcortical segmentation applied to isolate the subcortical grayordinate portion of the CIFTI space.

After the HCP minimal preprocessing pipelines, movement scrubbing was performed (85). As done in our prior work (86), all BOLD image frames with possible movement-induced artifactual fluctuations in intensity were flagged using two criteria: frame displacement (the sum of the displacement across all six rigid body movement correction parameters) exceeding 0.5 mm (assuming 50 mm cortical sphere radius) and/or the normalized root mean square (RMS) (calculated as the RMS of differences in intensity between the current and preceding frame, computed across all voxels and divided by the mean intensity) exceeding 1.6 times the median across scans. Any frame that met one or both of these criteria, as well as the frame immediately preceding and immediately following, were discarded from further preprocessing and analyses. Subjects with more than 50% frames flagged using these criteria were excluded. Next, a high-pass filter (threshold 0.008 Hz) was applied to the BOLD data to remove low frequency signals due to scanner drift. In-house Matlab code was used to calculate the average variation in BOLD signal in the ventricles, deep white matter, and across the whole grey matter, as well as movement parameters. These signals, as well as their first derivatives to account for delayed effects, were then regressed out of the grey matter BOLD time series as nuisance variables (as any change in the BOLD signal due to these variables are persistent and likely not reflecting neural activity) (87). It should be noted that using global signal regression to remove spatially persistent artifact is highly controversial in neuroimaging (88, 89), but it remains the field-wide gold standard (though see other recent and emerging approaches at (90, 91)).

Behavioral Symptom and Cognitive Data. The behavioral measures analyzed included the PANSS, an assessment of

psychosis symptom severity (27), and the Brief Assessment of Cognition in Schizophrenia (BACS) battery, which provided an assessment of cognitive functioning (92). Only control subjects with complete BACS measures were used for analyses; PANSS symptom scores were imputed for control subjects for whom the PANSS had not been administered under the assumption that these subjects were asymptomatic. Only patient subjects with complete PANSS and BACS measures were used in analyses. The BACS scores used here are presented as standardized Z-scores normalized to the mean and standard deviation of the control group for each BSNIP site, as done in prior work (11). The full PANSS battery is conventionally divided into three sub-scales: Positive symptom scale (7 items), Negative symptom scale (7 items), and General Psychopathology symptom scale (16 items). The BACS battery consists of 6 individual sub-scores (92). In total, this yielded 36 symptom variables per participant. Effects of symptom variation across assessment sites have been rigorously characterized in prior publications (93). Nevertheless, we explicitly tested for site effects in our analyses, described in detail below.

Principal Component Analysis of Behavioral Symptoms and Cross-validation. The principal component analysis (PCA) of behavioral data was computed by including all 36 symptom variables across all N=436 patients. Variables were first scaled to have unit variance across patients before running the PCA. Significance of the derived principal components (PCs) was computed via permutation testing. For each permutation, patient order was randomly shuffled for each symptom variable before re-computing PCA. This permutation was repeated 5,000 times to establish the null model. PCs which accounted for a proportion of variance that exceeded chance ($p < 0.05$ across all 5000 permutations) were retained for further analysis.

To evaluate if there were site differences that uniquely drove the PCA solution, we performed a leave-site-out cross-validation analysis. Specifically, we re-ran the PCA using all patients except those from one site, which was held out. This process was repeated for each of the 6 sites. To further evaluate the stability of the derived PCA solutions, we performed 1,000 runs of k -fold cross-validation for values of k between 2 and 10. For each k -fold run, the full sample of patients was randomly divided into equally-sized k sets and a PCA was re-computed using subjects in $k-1$ sets (as the left-out set was used as the held-out sample).

For each run of leave-site-out and k -fold cross-validations significance was assessed via permutation testing as described above. The number of significant PCs and the total proportion of variance explained by these significant PCs remained highly stable across all runs (see Fig. S4A and Fig. S6 and Fig. S5). Additionally, we compared observed and predicted PCA scores for the held-out sample of patients. Predicted scores for the held-out sample of patients were computed by weighing their raw symptom scores with the loadings from the PCA computed in all other patients. Observed scores for held-out patients were obtained from the original PCA computed on the full sample presented in the

main text. The similarity between predicted and observed scores was high for all five significant PCs across all runs of leave-site-out and k -fold cross-validation, exceeding $r=0.9$ in most analyses (see **Fig. S4B-C** and **Fig. S6** and **S5**). Notably, the PCA solution did not show medication status or dosage effects (**Fig. S7**).

We also assessed the similarity of the PCA loadings using leave-site-out and 1,000 runs of 5-fold cross-validation frameworks (**Fig. S4D**). Importantly, this cross-validation was designed to test if the observed loadings remained stable (as opposed to predicted patient-level scores). The loadings for significant PCs from each leave-site-out PCA solution as well as each run of the 5-fold cross-validation were highly correlated (**Fig. S4D**).

The leave-site-out and k -fold cross-validation PCA analyses by definition use overlapping samples of patients for each iteration. Therefore, we additionally conducted a full split-half replication using entirely non-overlapping sets of patients in each iteration. For each split-half iteration, the full patient sample was randomly divided into two sets with equal proportions of each of the three diagnostic groups (BPP, SADP, SZP). Then, a PCA was computed using each of the split-half patient samples. The loadings from the two PCA solutions were then evaluated for reproducibility. This process was repeated 1,000 times. The loadings for significant PCs were again highly similar even when comparing PCA solutions derived from completely non-overlapping patient samples (**Fig. S4D**).

To predict individual patient PC scores for the leave-one-out analysis (**Fig. 6A**), a PCA was computed using all patients except one held-out patient ($N=435$). In turn, the derived loadings were then used to compute the predicted PC scores for the left-out patient. This process was repeated until predicted PC scores were calculated for every patient. Finally, the predicted score for each patient was evaluated for reproducibility relative to the observed score obtained from the PCA solution computed using the full $N=436$ sample of patients. In addition, we computed an independent component analysis (ICA) to evaluate the consistency of the behavioral data reduction solution across methods (see).

Global Brain Connectivity Calculation. Following preprocessing, the functional connectivity (FC) matrix was calculated for each participant by computing the Pearson's correlation between every grayordinate in the brain with all other grayordinates. A Fisher's r -to- Z transform was then applied. Global brain connectivity (GBC) was calculate by computing every grayordinate's mean FC strength with all other grayordinates (i.e. the mean, per row, across all columns of the FC matrix). GBC is a data-driven summary measure of connectedness that is unbiased with regards to the location of a possible alteration in connectivity (94) and is therefore a principled way for reducing the number of neural features while assessing neural variation across the entire brain.

$$GBC(x) = \frac{1}{N} \sum_{y=1}^N r_{xy} \quad (1)$$

- where $GBC(x)$ denotes the GBC value at grayordinate x ;
- where N denotes the total number of grayordinates;
- where $\sum_{y=1}^N$ denotes the sum from $y = 1$ to $y = N$;
- where r_{xy} denotes the correlation between the time-series of grayordinates x and y ;

For parcel-wise GBC maps (described below) we first computed the mean BOLD signal within each parcel (see section below for parcellation details) for each participant and then computed the pairwise FC between all parcels. Finally, to obtain the parcellated GBC metric we computed the mean FC for each parcel and all other parcels. This order of operations (first parcellating the dense data series and then computing GBC) was chosen because it resulted in stronger statistical values due to increased within-parcel signal-to-noise of the BOLD data (**Fig. 3A**).

$$GBC(p) = \frac{1}{N} \sum_{q=1}^N r_{pq} \quad (2)$$

- where $GBC(p)$ denotes the GBC value at parcel p ;
- where N denotes the total number of parcels;
- where $\sum_{q=1}^N$ denotes the sum from $q = 1$ to $q = N$;
- where r_{pq} denotes the correlation between the time-series of parcels p and q ;

Neural Data Reduction via Functional Brain-wide Parcellation. Here we applied a recently developed Cole-Anticevic Brain Network Parcellation (CAB-NP) parcellation (33), which defines functional networks and regions across cortex and subcortex that leveraged the Human Connectome Project's Multi-Modal Parcellation (MMP1.0) (33, 36). The final published CAB-NP 1.0 parcellation solution can be visualized via the Brain Analysis Library of Spatial maps and Atlases (BALSA) resource (<https://balsa.wustl.edu/rrg5v>) and downloaded from the public repository (<https://github.com/ColeLab/ColeAnticevicNetPartition>). The cortex component of the parcellation solution is comprised of 180 bilateral parcels (a total of 360 across both left and right hemispheres), consistent with the Human Connectome Project's Multi-Modal Parcellation (MMP1.0) (36). The subcortex component is comprised of 358 parcels defined using resting-state functional BOLD covariation with the cortical network solution (33).

Mass Univariate Behavioral-to-Neural Mapping. Behavioral scores (*a priori* and PCA) were quantified in relation to individual GBC variation (either dense or parcellated) via a mass univariate regression procedure. The resulting map of

regression coefficients reflected the strength of the relationship between patients' behavioral PC score and GBC at every neural location, across all 436 patients. The greater the magnitude of the coefficient for a given location, the stronger the statistical relationship between GBC and the behavioral variation across patients. The coefficients were then Z-scored for each map. Significance of the maps was assessed via non-parametric permutation testing, 2000 random shuffles with TFCE (95) type-I error-protection computed via the Permutation Analysis of Linear Models program (96).

Cross-validation of Mass Univariate Behavioral-to-Neural Mapping. Behavioral-to-Neural mapping followed the same cross-validation logic as described for the symptom-driven PCA solutions. Specifically, five-fold cross-validation was performed by first randomly partitioning all patients (N=436) into 5 subsets. Regression of the behavioral PC scores onto GBC across patients was performed while holding out 1/5 of the patient sample (N=349). The correlation between the resulting neural coefficient map was then computed with the neural map obtained from the full sample calculation.

For leave-site-out cross-validation, all subjects except those from one site were used when calculating the Behavioral-to-Neural regression model. The resulting maps were compared to the map from the full sample regression. This process was repeated 6 times, each time leaving out subjects in one of the 6 sites.

Additionally, 1,000 iterations of split-half replication of the neural-behavioral mapping were performed. For each split-half replication iteration, the full sample of subjects was first randomly split into two halves (referred to as H1 and H2) with the proportion of subjects in each diagnostic group (BPP, SADP, SZP) preserved within each half. For each iteration we used the H1 PCA solution and loadings to compute the predicted PCA scores for H2. In turn, the observed H1 scores were computed from a PCA loadings on the same H1 half-sample of patients. These H1 scores were then regressed against parcellated GBC for patients in H2. This coefficient GBC map reflects the strength of the relationship between the predicted PC score and GBC across H2 patients. Finally, the GBC coefficient maps derived from the H1 observed and H2 predicted PCA scores were then correlated for each PC axis. This process was then repeated 1,000 times and evaluated for reproducibility.

Principal Component Analysis of Neural Data. To evaluate the consistency of the neural GBC, we computed a PCA solution for the parcellated neural GBC data for all 202 control subjects as well as for all 436 patients (Fig. S17). The resulting neural GBC-derived PCs capture the striking consistency of the neural variance components irrespective of clinical status, which highlights that the bulk of the neural variance is not symptom-relevant. As with the behavioral PCA, significance of the neural PCA solution was assessed via permutation testing (1,000 random shuffles of parcels within subject).

Canonical Correlation Analysis. Canonical correlation analysis (CCA) is a multivariate statistical technique which

examines simultaneously the relationships between multiple independent variables and multiple dependent variables by computing linear combinations of each variable set that maximizes the correlations between the two sets (Fig. 4A). Here, the two variates used were behavioral features across all subjects and neural features across all subjects. Here, each feature was Z-scored prior to computing the CCA solution. Given the size of the 'dense' neural feature space reduced the number of neural features in principled manner via the described parcellation. This reduced the number of neural features to 180 symmetrized cortical parcels (i.e. GBC was averaged for each pair of analogous parcels in the left and right hemispheres). Critically, corresponding cortical parcels in the left and right hemispheres of this parcellation have been shown to be highly similar ((31, 81)). First, we computed the CCA solution using all 36 behavioral symptom item-level measures as behavioral features. In turn, we computed an additional CCA solution using 5 significant PC scores.

Each behavioral PC is a weighted linear composite of the original behavioral items and each behavioral CV is a weighted linear composite of the behavioral PCs. Therefore, to compute the loadings of the 36 original behavioral items on each of the behavioral CVs (computed using the behavioral PCs), we multiplied the matrix of loadings from the CCA with the matrix of loadings from the PCA (Fig. 4J).

In addition to the CCA computed between 180 neural cortical parcel GBC features and 36 behavioral items or 5 PC scores (shown in the main text, Fig. 4), we also computed the following control CCAs: i) 5 PCs and 358 subcortical parcel GBC; ii) 5 PCs and GBC from 12 whole-brain networks. Notably, both the 358 subcortical parcels and the 12 functional networks (which span both cortex and subcortex) were obtained from the parcellation in (33).

CCA Cross-validation. Five-fold cross-validation of the CCA was performed by first randomly partitioning all patients (N=436) into 5 subsets. The CCA was then performed between neural and behavioral features using all but one of the subsets. The results of this CCA was then compared to the full sample CCA. Fig. S16A-D shows the comparisons across all 5 five-fold cross-validation runs of the CCA compared to the full model, including neural factor loadings, behavioral factor loadings, and projected behavioral item loadings.

For leave-site-out cross-validation, all subjects except those from one site were used in the CCA. The resulting outputs were then compared to those from the full sample CCA. This process was repeated for all 6 sites. These data are shown in Fig. S16E-H.

Split-half replication of the CCA was performed by first randomly split into two halves (referred to as H1 and H2) with the proportion of subjects in each diagnostic group (BPP, SADP, SZP) preserved within each half. A CCA was then performed separately in each half and the resulting outputs were then compared to each other. This process was repeated 1,000 times to obtain mean and standard deviation performance metrics. These data are shown in Fig. 4L.

For leave-one-out cross-validation of the CCA, one subject

was held out as a CCA was performed using neural and behavioral features from the other 435 subjects. The loadings matrices Ψ and Θ from the CCA were then used to calculate the “predicted” neural and behavioral latent scores for all 5 canonical modes for the holdout subject. This was repeated for every subject, such that predicted neural and behavioral latent score matrices (\hat{N} and \hat{B}) were computed, of the same dimensions as (N and B) respectively. The corresponding CVs in \hat{N} and \hat{B} were then correlated across subjects, as shown in **Fig. 4M**. If the CCA solution is stable, these correlations should be comparable to the canonical correlations of the full CCA (**Fig. 4D**).

Independent Replication Dataset. To illustrate the generalizability of the single-subject prediction results across independent datasets and DSM disorders, we use an independent collected dataset consisting of 30 patients with a formal diagnosis of schizophrenia (SZP) and 39 patients diagnosed with obsessive compulsive disorder (OCD). These patients were recruited via clinician referral and regional advertising and assessed at the Connecticut Mental Health Center in New Haven, CT. Demographic and symptom data are shown in **Table S3**.

The behavioral assessment for the replication dataset included the PANSS and the Penn Computerized Neurocognitive Battery (PennCNB)(97). Items from the PennCNB were matched to items from the BACS as follows: BACS digital sequencing (working memory) -> LNB (N-back Total score); BACS symbol coding (attention) -> Penn Continuous Performance Test); BACS token motor task (motor) -> CTAP; BACS Tower of London (abstraction and mental flexibility) -> PCET or RAVEN (Penn Conditional Exclusion Test); BACS verbal fluency (verbal ability) -> SPVRT (Penn Verbal Reasoning Test); BACS verbal memory test -> PLLT (Penn List Learning Test) (98).

Neural data were collected using a Siemens 3T scanner with a 64 channel head coil at the Yale Center for Biomedical Imaging. Imaging acquisition parameters were aligned with those of the Human Connectome Project (HCP) (99). High-resolution T1w and T2w structural images were acquired in 224 AC-PC aligned slices, 0.8mm isotropic voxels. T1w images were collected with a magnetization-prepared rapid gradient-echo (MP-RAGE) pulse sequence (TR=2400 ms, TE=2.07 ms, flip angle=80, field of view=256 x 256 mm). T2w images were collected with a SCP pulse sequence [TR=3200 ms, TE=564 ms, flip angle mode=T2 var, field of view=256 x 256 mm]. Resting-state BOLD images were collected with a multi-band accelerated fast gradient-echo, echo-planar sequence (acceleration factor=6, time repetition (TR)=800 ms, time echo (TE)=31.0 ms, flip angle=550, field of view=210 x 210 mm, matrix=84 x 84, bandwidth=2290 Hz); 54 interleaved axial slices aligned to the anterior-posterior commissure (AC-PC) with 2.5mm isotropic voxels. Additionally, a pair of reverse phase-encoded spin-echo field maps (anterior-to-posterior and posterior-to-anterior) were acquired (voxel size=2.5 mm isotropic, TR=7220 ms, TE=73 ms, flip angle=900, field of view=210 x 210 mm, bandwidth=2290 Hz). These neural data were preprocessed using

the HCP minimal preprocessing pipeline as described in the section “Neural data acquisition and preprocessing” above, with the same motion scrubbing and nuisance signal regression parameters as were used for the BSNIP dataset.

Neural Parcel ΔGBC Feature Selection and Prediction.

For the purposes of patient selection, we were focused on individual differences in neural features as they co-vary in relation to individual symptom scores. However, this individual variation in neural features may be small compared to the overall group mean. Hence for each patient, we compute the difference in each brain location (parcel) between the patient’s actual GBC and the group mean GBC for that location. This is denoted by ΔGBC . Importantly, using a de-meaned GBC metric “standardizes” the data and helps to correct for possible differences in scanners/protocols across different datasets. Next, we developed an optimized univariate regression framework leveraging a “dot product” metric to relate a vector of neural features with a single symptom scalar value. This process for neural feature selection (results in **Fig. 6**) is shown as a systems flow diagram in **Fig. S18**. The observed dot product GBC metric ($dpGBC^{obs}$) is computed as follows:

$$dpGBC_i^{obs} = \begin{pmatrix} \Delta GBC_1^{obs} \\ \vdots \\ \Delta GBC_{P_{select}}^{obs} \end{pmatrix}_i \cdot \begin{pmatrix} \beta GBC_1^{obs} \\ \vdots \\ \beta GBC_{P_{select}}^{obs} \end{pmatrix}_{N-i} \quad (3)$$

- where for a given subject i $dpGBC_i^{obs}$ denotes the dot product GBC value of the two vectors ΔGBC_i^{obs} and βGBC_{N-i}^{obs} across all P_{select} parcels,
- ΔGBC_i^{obs} is a vector of length P_{select} denoting a difference map of subject i ’s GBC^{obs} map relative to the group mean GBC^{obs} map, within a given number of parcels P_{select} ,
- the βGBC_{N-i}^{obs} vector denotes the PCA-to-GBC statistical group-level β map for a low-dimensional PC symptom score across selected parcels P_{select} for N subjects excluding subject i .

This calculation is then repeated for each subject i , resulting in a final vector $dpGBC^{obs} = [dpGBC_1^{obs}, \dots, dpGBC_N^{obs}]$ for N subjects. There are several key properties of the $dpGBC^{obs}$ statistic: i) it is not inflated by individual GBC map similarity to the group map because each subject’s ΔGBC^{obs} map is demeaned relative to the reference group computed independently of the left-out subject; ii) this statistic is not biased by the parcel number (which drops with iterative selection) because the resulting $dpGBC^{obs}$ value variation is quantified relative to the low-dimensional symptom score across all subjects (see Eq. 2). Put differently, the final evaluation considers the relationship between $dpGBC^{obs}$ and the low-dimensional PC symptom scores across individuals; iii) The dot product statistic can yield both positive and negative values – a property which some map similarity measures lack (e.g. η^2); iv) It is unbounded (unlike a

correlation), which is key to maximize co-variation with low-dimensional symptom scores across individuals (see Eq. 2); v) The ΔGBC_i^{obs} map for a given individual is projected onto the basis set of the βGBC_{N-i}^{obs} map, which is independent of the left-out individual but directly related to the low-dimensional PC symptom score variance, thus maximizing the dot product optimization.

Next, we select $N - i$ individuals and compute a univariate regression where participants' low dimensional symptom scores are regressed onto the $dpGBC_{N-i}^{obs}$ values:

$$\begin{pmatrix} dpGBC_1^{obs} \\ \vdots \\ dpGBC_{N-i}^{obs} \end{pmatrix} = \alpha \begin{pmatrix} S_1^{obs} \\ \vdots \\ S_{N-i}^{obs} \end{pmatrix} + \epsilon \quad (4)$$

- where for $dpGBC^{obs}$ each element denotes the dot product value of the $\Delta GBC_i^{obs} \cdot \beta GBC_{N-1}^{obs}$ vectors per subject across P parcels,
- α denotes the regression coefficient in the univariate linear model,
- where for S^{obs} each element denotes the observed low-dimensional symptom score (e.g. $PC3$ score) per subject,
- ϵ denotes the error term in the univariate linear model.

After the $dpGBC^{obs} = \alpha S^{obs} + \epsilon$ regression is computed on $N - i$ subjects it is applied to the left-out-subject i . This is repeated for all N subjects and the model is evaluated for the number of parcels that maximize two key dot product evaluation metrics (**Fig. 5B**):

Metric A

For $p = \{1, \dots, P\}$ find p where

$$\Rightarrow A = \max_r \left[\begin{pmatrix} dpGBC_1 \\ \vdots \\ dpGBC_N \end{pmatrix}^{pred}, \begin{pmatrix} S_1 \\ \vdots \\ S_N \end{pmatrix}^{pred} \right] \quad (5)$$

Metric B

For $p = \{1, \dots, P\}$ find p where

$$\Rightarrow B = \max_r \left[\begin{pmatrix} dpGBC_1 \\ \vdots \\ dpGBC_N \end{pmatrix}^{obs}, \begin{pmatrix} dpGBC_1 \\ \vdots \\ dpGBC_N \end{pmatrix}^{pred} \right] \quad (6)$$

- where A denotes the maximum r correlation value for the two vectors of $dpGBC^{pred}$ values and the predicted S^{pred} low-dimensional symptom scores (e.g. $PC3$ axis) for N subjects from the leave-one-out cross-validation,
- where B denotes the maximum r correlation value for the two vectors of obs and $pred$ $dpGBC$ values for N subjects.

In the initial step in the step-down model all $P = 718$ parcels are retained in the initial dot product calculation. For each iteration of P selected parcels, the least predictable parcel P (i.e. the parcel with the weakest value in the $PC3$ map) is eliminated from the map. Then, the step-down regression is repeated until $P = 1$.

Pharmacological Neuroimaging Acquisition in Healthy Volunteers - LSD.

Methods for the lysergic acid diethylamide (LSD) neuroimaging study are described in detail in prior publications (28). The study employed a fully double-blind, randomized, within-subject cross-over design with 3 conditions: (1) placebo + placebo (Pla) condition: placebo (179 mg Mannitol and Aerosil 1 mg po) after pretreatment with placebo (179 mg Mannitol and Aerosil 1 mg po); (2) Pla+LSD (LSD) condition: LSD (100 μ g po) after pretreatment with placebo (179 mg Mannitol and Aerosil 1 mg po), or (3) Ketanserin+LSD (Ket+LSD) condition: LSD (100 μ g po) after pretreatment with the 5-HT_{2A} antagonist Ket (40 mg po). Data were collected for h subjects in a randomized counterbalanced order at three different sessions each two weeks apart. For all conditions, the first substance was administered 60 minutes before the second substance, and the first neural scan was conducted 75 minutes after the second administration, with a second scan conducted at 300 minutes post-administration. In the present study, only data from the two neural scans for the LSD and Pla conditions were evaluated.

Briefly, neuroimaging data acquisition details for the LSD study are as follows. MRI data were acquired on a Philips Achieva 3.0T whole-body scanner (Best, The Netherlands). A 32-channel receiver head coil and MultiTransmit parallel radio frequency transmission was used. Images were acquired using a whole-brain gradient-echo planar imaging (EPI) sequence (repetition time=2,500 ms; echo time=27 ms; slice thickness=3 mm; 45 axial slices; no slice gap; field of view=240 mm²; in-plane resolution=3 mm \times 3 mm; sensitivity- encoding reduction factor=2.0). 240 volumes were acquired per resting state scan resulting in a scan duration of 10 mins. Additionally, two high-resolution anatomical images were acquired using T1-weighted and T2-weighted sequences. T1-weighted images were collected via a 3D magnetization-prepared rapid gradient-echo sequence (MP-RAGE) with the following parameters: voxel size=0.7 mm³, time between two inversion pulses=3123 ms, inversion time=1055 ms, inter-echo delay=12 ms, flip angle=8°, matrix=320 \times 335, field of view=224 \times 235 mm, 236 sagittal slices. Furthermore T2-weighted images were collected using via a turbo spin-echo sequence with the following parameters: voxel size=0.7 mm³, repetition time=2500 ms, echo time=415 ms, flip angle=90°, matrix=320 \times 335, field of view=224 mm \times 235 mm, 236 sagittal slices.

Pharmacological Neuroimaging Acquisition in Healthy Volunteers - Ketamine.

Similar to the LSD study, the ketamine pharmacological neuroimaging protocol employed a within-subject design where all healthy volunteer participants underwent a single scanning session consisting of two infusions:

i) placebo (saline solution) followed by ii) ketamine infusion. Healthy volunteers were informed prior to scanning that they would undergo one placebo run and one ketamine run but were blinded to the order of administration. Because of the sustained effects of ketamine, this infusion was always the second of the two runs, consistent with prior work (86). A doctor and two nurses as well as a study coordinator remained present for the duration of the scan. A bolus of ketamine (0.3 mg/kg of bodyweight) or saline were delivered via infusion 5 sec after the start of the run and then continuously at a rate of 0.65 mg/kg through the duration of the session. The sequence of scans in each for either the placebo or ketamine infusion was as follows: i) resting state (4.67 min); ii) blood draw (sham if saline condition); iii) a cognitive working memory task (total 14 min); iv) blood draw (sham if saline condition); v) a cognitive working memory task (total 14 min); vii) blood draw (sham if saline condition); vii) a cognitive working memory task (total 8.63 min). Data from the the cognitive working memory task were not used in the present study and are actively undergoing a distinct analysis. Participants were scanned at Yale University on a Siemens Trio 3T whole-body scanner and 32-channel receiver head coil. High-resolution structural T1-weighted images were acquired using an MP-RAGE sequence and the following parameters: voxel size=0.8 mm³, time between two inversion pulses=3123 ms, inversion time=1055 ms, inter-echo delay=12 ms, flip angle=8°, matrix=320×335, field of view=227×272 mm, 227 sagittal slices. T2-weighted images were acquired with the following parameters: voxel size=0.8 mm³, repetition time=2500 ms, echo time=415 ms, flip angle=90°, matrix=320 × 335, field of view=227×272 mm, 227 sagittal slices. BOLD images were acquired using a whole-brain gradient-echo planar imaging (EPI) sequence (400 frames at TR=0.7 ms; TE=30 ms; slice thickness=2.5 mm; 54 axial slices; no slice gap; field of view=250 mm²; in-plane resolution=2.5 mm × 2.5 mm). In addition, a field map and pair of spin-echo gradients were collected at the end of every scanning session.

Pharmacological Neuroimaging Processing and Analysis for LSD and Ketamine Samples. All data were preprocessed to be consistent with the BSNIP and replication dataset processing steps. Specifically, we used the HCP minimal pre-processing pipeline and with the same motion scrubbing and nuisance signal regression parameters as were used for the BSNIP dataset as described in the section “Neural data acquisition and preprocessing” above. Parcel-wise GBC maps were computed as described in the section “Resting-state functional connectivity and global brain connectivity” above, by first parcellating the dense time-series for each subject and then computing the parcel-level GBC. As with the BSNIP and replication analyses, parcel GBC was first calculated by parcellating the dense time-series for every subject and then computing GBC across all 718 parcels from the whole-brain functional parcellation (33, 36). Group-level GBC maps were computed for every participant for all conditions. We computed the contrast map for “LSD-Placebo” as well as “Ketamine-Placebo” conditions as a t-test between both the

pharmacological scans versus placebo scans. The Z-scored t-contrast maps between pharmacological and placebo conditions were used as a pharmacological target map in relation to the PC-derived neural GBC maps **Fig. 7** and **Fig. S22**.

Neural Gene Expression Mapping. The gene mapping analyses in this study utilize the procedure described in (31). Briefly, we used cortical gene expression data from the publicly available Allen Human Brain Atlas (AHBA, RRID:SCR_007416), mapped to cortex (31). Specifically, the AHBA quantified expression levels across 20,737 genes obtained from six postmortem human brains using DNA microarray probes sampled from hundreds of neuroanatomical loci. Recent studies demonstrated the ability to map expression of each gene onto neuroimaging-compatible templates (30, 31). Building on these innovations we mapped gene expression on to 180 symmetrized cortical parcels from the HCP atlas (36) in line with recently published methods (31). This yielded a group-level map for each gene where the value in each parcel reflected the average expression level of that gene in the AHBA dataset. These group-level maps were in turn as a gene expression target map in relation to the PC-derived neural GBC maps.

Data Availability. All primary results derive from data that is publicly available from sources described above. All symptom and neuroimaging data for the BSNIP-1 consortium was obtained via the NDA (https://nda.nih.gov/edit_collection.html?id=2274).