RESEARCH ARTICLE

NEURODEVELOPMENT

Contrastive machine learning reveals the structure of neuroanatomical variation within autism

Aidas Aglinskas*, Joshua K. Hartshorne, Stefano Anzellotti

Autism spectrum disorder (ASD) is highly heterogeneous. Identifying systematic individual differences in neuroanatomy could inform diagnosis and personalized interventions. The challenge is that these differences are entangled with variation because of other causes: individual differences unrelated to ASD and measurement artifacts. We used contrastive deep learning to disentangle ASD-specific neuroanatomical variation from variation shared with typical control participants. ASD-specific variation correlated with individual differences in symptoms. The structure of this ASD-specific variation also addresses a long-standing debate about the nature of ASD: At least in terms of neuroanatomy, individuals do not cluster into distinct subtypes; instead, they are organized along continuous dimensions that affect distinct sets of regions.

Psychiatric disorders affect millions of people worldwide. Heterogeneity is a major obstacle to understanding them: Individuals diagnosed with the same disorder often present with different behavioral symptoms and genetic variants (*1*). We investigated heterogeneity within autism spectrum disorder (ASD), a prevalent neurodevelopmental condition (*2*) characterized by impaired social interactions, restricted patterns of behavior, and communication deficits (*3*). Individuals with ASD differ in the severity of behavioral symptoms (*4*), in their genetics (*5*), and in neuroanatomy (*6*).

Understanding neuroanatomical heterogeneity within ASD could be pivotal to improving quality of life in affected individuals, by leading to more specific diagnoses and targeted behavioral interventions (*7*, *8*). However, researchers have not yet identified systematic neuroanatomical variation that correlates with symptoms and that generalizes across different groups of participants (*6*).

We hypothesized that ASD-specific variation has been obscured by other factors that lead brains to vary. Brains differ from one to another because of numerous genetic and environmental causes unrelated to ASD (*9*). Neuroanatomical data from different individuals also varies because of methodological artifacts, such as systematic differences between scanners and scanning sites (*10*). ASD-specific variation may be difficult to identify within this mass of irrelevant variation. Methods in use now for addressing these problems remain unsatisfactory. For instance, matching ASD and typical control (TC) participants works in theory, but it assumes that we know which factors we need to match. However, brain anatomy is shaped by a multitude of genetic and environmental factors (*9*), some of which are unknown, undermining any attempt at matching.

To better characterize ASD-specific neuroanatomical variation, we disentangled it from variation that is common to the general population using contrastive variational autoencoders (CVAEs) (*11*, *12*). CVAEs take as inputs samples from two distinct populations and isolate variation specific to one population from variation common to both (fig. S1). We used CVAEs to disentangle "ASD-specific" neuroanatomical variation from variation "shared" by both ASD and TC participants, representing each as a distinct set of latent features (Fig. 1A). First, we validated the features by confirming that the ASD-specific features are differentially related to clinical symptoms, whereas the shared features are differentially related nonclinical properties. We replicated the results with a zero-free-parameter generalization to an independent dataset. Next, we applied cluster analysis to the ASD-specific features to determine whether there are distinct subtypes of ASD neuroanatomy. Finally, we leveraged the properties of the CVAE to identify brain regions that vary systematically within the ASD population.

Results

ASD-specific neuroanatomy relates to clinical variation

We used the Autism Brain Imaging Data Exchange I (ABIDE I) magnetic resonance imaging (MRI) dataset [(*13*); 470 ASD participants, 512 TCs] to train a CVAE and a noncontrastive VAE that has a single set of latent features but is matched to the CVAE in the number of parameters and in the number of latent features. The noncontrastive VAE allows us to test whether associations between neuroanatomy and ASD symptoms can be identified using variational autoencoding alone, without disentangling ASD-specific and shared variation.

Thus, to establish a baseline, we first report the noncontrastive VAE results. We used representational similarity analysis (RSA) (*14*) to test whether the VAE's neuroanatomical features correlate with individual variation in the ASD participants' nonclinical and clinical characteristics, such as scanner type, age, Vineland adaptive behavior scores, and Autism Diagnostic Observation Schedule (ADOS) scores (a numerical measure of ASD symptom severity). We first calculated the pairwise dissimilarity between participants with respect to the VAE neuroanatomical features and obtained a dissimilarity matrix. We then repeated this process for each nonclinical and clinical characteristic (Fig. 1B). Finally, we compared the VAE dissimilarity matrix to the matrices for each individual characteristic using the Kendall rank correlation coefficient (Kendall τ).

The VAE features showed Kendall τ correlations with some of the nonclinical characteristics, such as scanner type ($\tau = 0.04$, $t_9 = 16.29$, $p < 0.001$), age ($\tau = 0.03$, $t_9 = 8.27$, $p < 0.001$), and gender ($\tau = 0.03$, $t_9 = 4.71$, $p = 0.001$). Whereas there was some relationship between neuroanatomical feature similarity extracted by VAE and Diagnostic Statistical Manual IV (DSM IV) behavioral subtypes ($\tau = 0.03$, $t_9 = 4.77$, $p = 0.001$), there was no relationship with autism severity (ADOS total; $\tau = 0.00$, $t_9 = -1.08$, $p = 0.310$) or Vineland adaptive behavior scores ($\tau = 0.00$, $t_9 = -0.29$, $p = 0.780$). This is consistent with the idea presented above that entangled measures of neuroanatomy (such as VAE features) may fail to capture variation in symptoms.

We then assessed whether disentangling ASD-specific and shared neuroanatomical variation with a CVAE would allow us to identify clinically relevant individual variation. As described above, the CVAE segregates its internal representations into ASD-specific and shared features (Fig. 1A and fig. S2). Although the CVAE training implicitly makes a binary distinction between ASD and TC participants, the model is not provided with any of the clinical and nonclinical individual characteristics of interest. We used RSA to compare the CVAE's ASD-specific and shared neuroanatomical features to each of the individual characteristics. We expected to find that shared features correlate with nonclinical variation that is common to both ASD and TC participants, whereas ASD-specific features correlate with clinical ASD variation (Fig. 1B).

As expected, scanner type was associated with subject similarity in the shared features ($\tau = 0.11$, $t_9 = 253.01$, $p < 0.001$) but not the ASD-specific features ($\tau = -0.01$, $t_9 = -14.16$, $p < 0.001$; shared versus ASD-specific: $\Delta\tau = 0.12$,

Department of Psychology and Neuroscience, Boston College, Boston, MA 02467, USA.
*Corresponding author. Email: aidas.aglinskas@gmail.com

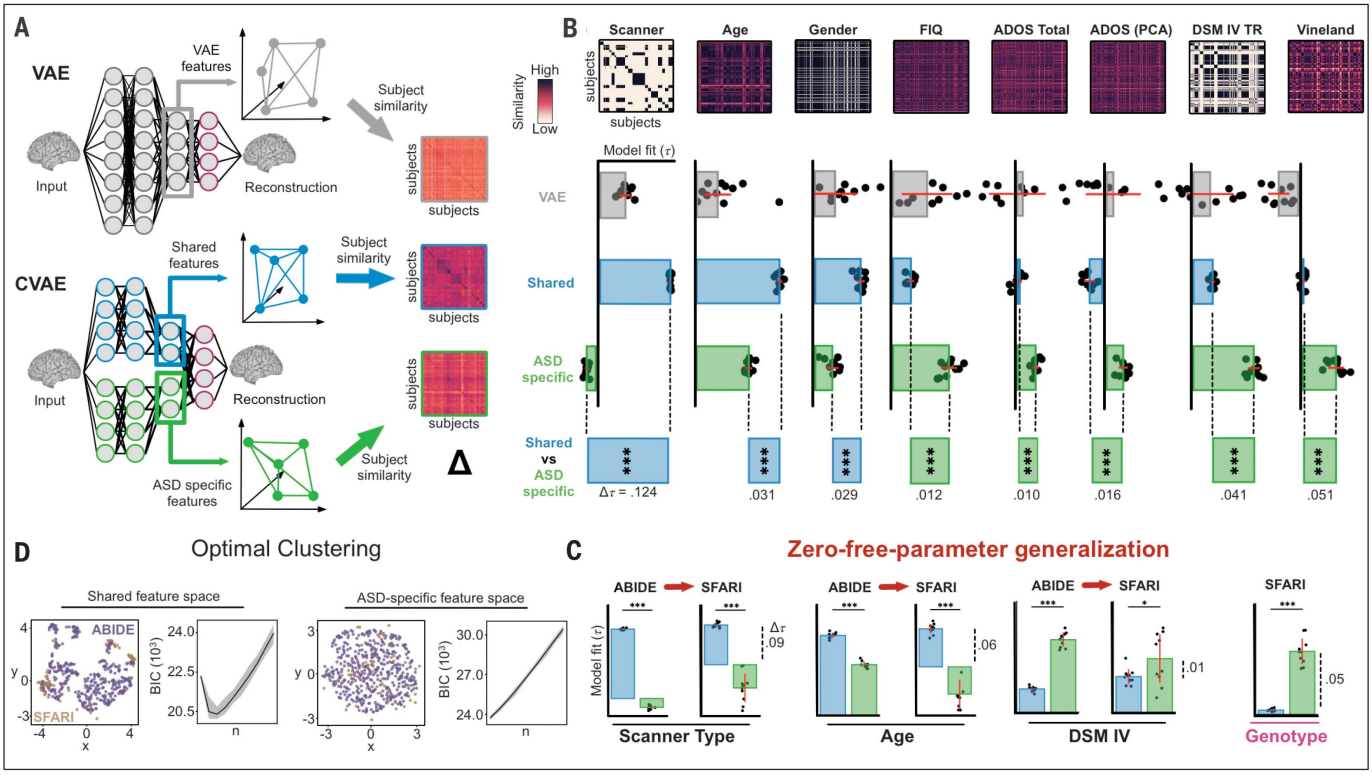

Fig. 1. Neuroanatomical feature models. (**A**) Neuroanatomical features extracted from the autoencoders are used to construct neuroanatomical similarity matrices. (**B**) Neuroanatomical similarity matrices are compared with similarity based on different participant properties. Variables common to TC and ASD participants are best captured by the shared CVAE features, and variables associated with ASD-related variation are best captured by the ASD-specific features. Model fit for the control model (VAE) is worse across all variables. Red horizontal lines 95% confidence intervals. PCA, principal components analysis; DSM IV TR, DSM IV Text Revision. (**C**) Zero-free-parameter generalization. The results generalize to a new dataset (SFARI) without the need for additional fitting; in addition, participants with the same CNV associated with increased risk of ASD (16p11.2 deletion or duplication) are more similar in ASD-specific, but not shared, neuroanatomical features. Red vertical lines indicate 95% confidence intervals. (**D**) Optimal clustering. Individual variation in ASD-specific features is best captured by a single cluster, whereas variation in the shared features is best captured by three clusters. Scatterplots show individual subjects' neuroanatomical data from ABIDE (purple) and SFARI (orange) datasets projected onto uniform manifold approximation and projection (UMAP) dimensions computed from the shared and ASD-specific features. *$p < 0.05$; ***$p < 0.0001$.

$t_9 = 124.83$, $p < 0.001$). Thus, the CVAE was able to factor out a common source of "nuisance" variation in multisite data (*10*). By contrast, measures of ASD clinical symptoms were more associated with the ASD-specific features but generally not associated with the shared features. These include DSM IV behavioral subtypes (ASD-specific: $\tau = 0.06$, $t_9 = 30.83$, $p < 0.001$; shared: $\tau = 0.02$, $t_9 = 29.02$, $p < 0.001$; comparison: $\Delta\tau = 0.04$, $t_9 = 20.04$, $p < 0.001$), ADOS total score (ASD-specific: $\tau = 0.01$, $t_9 = 16.85$, $p < 0.001$; shared: $\tau = 0.00$, $t_9 = -1.50$, $p = 0.167$; comparison: $\Delta\tau = 0.01$, $t_9 = 11.59$, $p < 0.001$), and Vineland adaptive behavior questionnaire (ASD-specific: $\tau = 0.05$, $t_9 = 12.33$, $p < 0.001$; shared: $\tau = 0.00$, $t_9 = 1.17$, $p = 0.270$; comparison: $\Delta\tau = 0.05$, $t_9 = 10.46$, $p < 0.001$) [see also fig. S4 and supplementary materials (SM)].

Results for age, gender, and full-scale intelligence quotient (FIQ) were of particular interest, because these are known to differently relate to neuroanatomy in the TC and ASD populations (*15*). Each of these properties was significantly related to both the ASD-specific features (age: $\tau = 0.05$, $t_9 = 48.60$, $p < 0.001$; gender: $\tau = 0.02$, $t_9 = 8.13$, $p < 0.001$; FIQ: $\tau = 0.02$, $t_9 = 20.22$, $p < 0.001$) and the shared features (age: $\tau = 0.08$, $t_9 = 89.29$, $p < 0.001$; gender: $\tau = 0.05$, $t_9 = 35.34$, $p < 0.001$; FIQ: $\tau = 0.01$, $t_9 = 15.57$, $p < 0.001$), suggesting that the CVAE was able to disentangle general effects of age, gender, and FIQ from those that specifically interact with ASD. Shared features captured greater variation in age and gender than ASD-specific features (age: $\Delta\tau = 0.03$, $t_9 = 24.11$, $p < 0.001$; gender: $\Delta\tau = 0.03$, $t_9 = 11.90$, $p < 0.001$). Conversely, variation in FIQ was more related to ASD-specific features than to shared features ($\Delta\tau = 0.01$, $t_9 = 12.86$, $p < 0.001$).

In sum, the CVAE was not only able to disentangle individual neuroanatomical variation that is specific to ASD from variation that characterizes the population as a whole, but these patterns of variation were differentially associated with clinical and nonclinical participant characteristics. This contrasts with the control VAE model, where unitary neuroanatomical features showed weaker correlations with individual characteristics.

Generalization to an independent dataset

Generalization to a new dataset is considered a gold-standard test of a model. Generalization across datasets is desirable, because a model trained on one group of participants may need to be used to inform the diagnosis of new participants that were not included in the training dataset. To test generalization, we applied the ABIDE-trained CVAE to the anatomical scans of participants from the Simons Foundation Autism Research Initiative (SFARI) Variation in Individuals Project (VIP) dataset ($N = 121$) (*16*) using a parameter-free fit (without retraining or transfer-learning).

Evaluating the performance of the model with this new dataset—collected by different researchers at different facilities—provides a more stringent test of generalization than does cross-validation [compare (*17*)]. However,

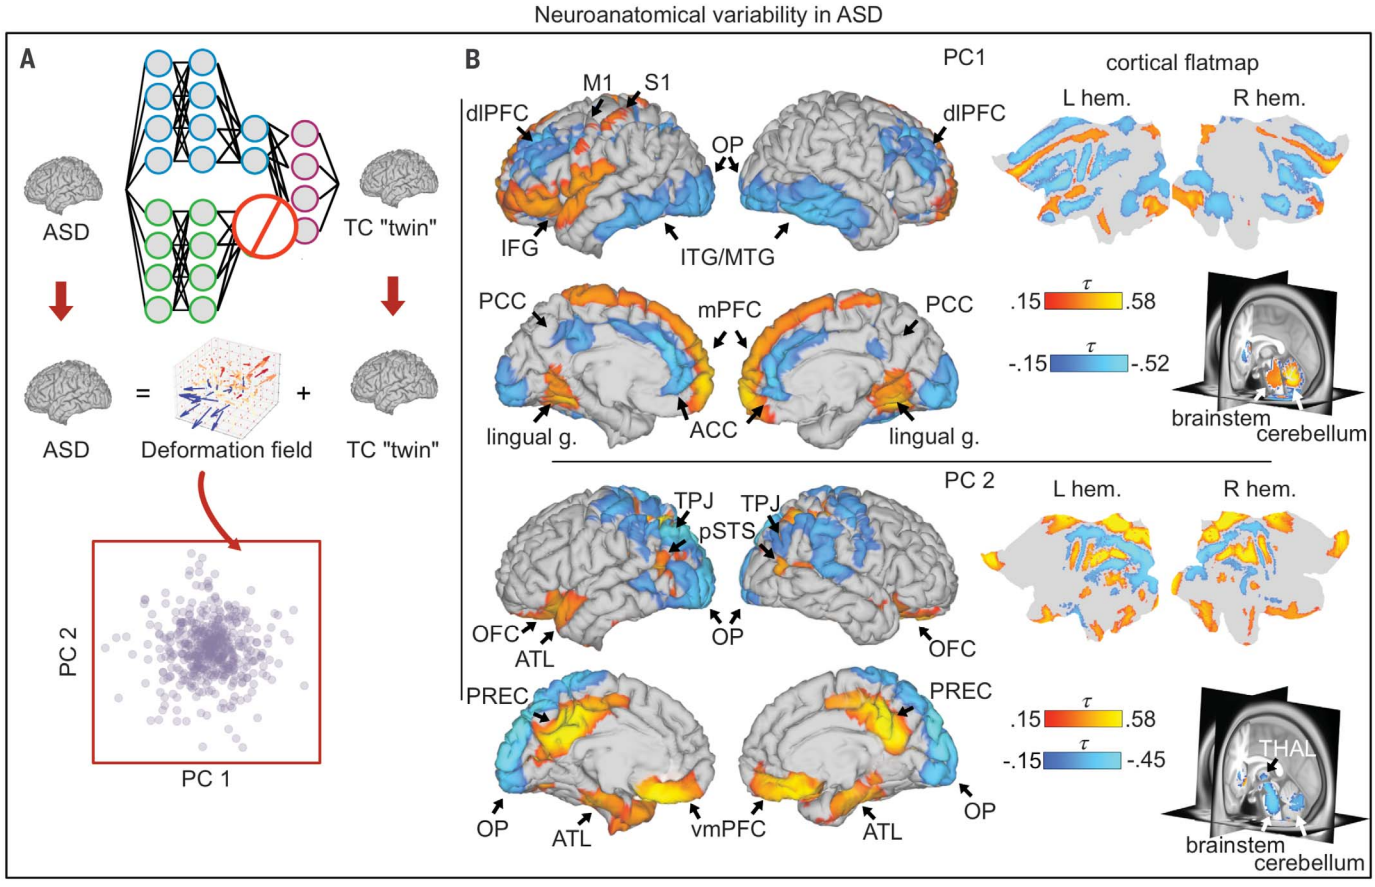

Fig. 2. Anatomical loci of individual variation within the ASD group. (**A**) For each ASD subject, we calculated a synthetic TC-twin brain matched on ASD-unrelated (shared) neuroanatomical features and morphed it into the corresponding ASD brain, obtaining a deformation field. We then applied principal components analysis to the Jacobian determinants of the deformation fields across participants. (**B**) Areas showing volumetric increases (red) and decreases (blue) associated with the two PCs that explain most variance. White matter effects are reported in fig. S8; analyses using diffusion weighted imaging will be needed to determine more precisely which specific tracts are affected. ACC, anterior cingulate cortex; ATL, anterior temporal lobe; dlPFC, dorsolateral prefrontal cortex; IFG, inferior frontal gyrus; ITG/MTG, inferior and middle temporal gyrus; L hem., left hemisphere; lingual g., lingual gyrus; M1, motor cortex; mPFC, medial prefrontal cortex; OFC, orbitofrontal cortex; OP, occipital pole; PCC, posterior cingulate cortex; PREC, precuneus; pSTS, posterior superior temporal sulcus; R hem., right hemisphere; S1, somatosensory cortex; THAL, thalamus; TPJ, temporoparietal junction; vmPFC, ventromedial prefrontal cortex.

ensuring that a machine-learning model will generalize to an arbitrary new dataset (or indeed, ensuring that any scientific finding will generalize) remains a difficult problem; more extensive testing of the generalizability of the results will require additional datasets.

SFARI VIP includes information about ASD-relevant copy number variations (CNVs), allowing us to study whether ASD-specific neuroanatomical features correlate with genotype. We performed analyses on the independent SFARI dataset that were identical to those performed on the ABIDE dataset, extracting shared and ASD-specific features and comparing neuroanatomical feature similarities in shared and ASD-specific features to subject properties similarities in scanner type, age, gender, DSM IV behavioral subtypes, and genotype.

We expected that if CVAE features were robust, then shared features should again differentially correlate with properties of scanning site, age, and gender, whereas ASD-specific features should correlate with ASD-related properties such as DSM IV subtypes. Results confirmed these predictions. Compared with ASD-specific features, shared features correlated better with scanner type ($\Delta\tau = 0.09$, $t_9 = 12.81$, $p < 0.001$), age ($\Delta\tau = 0.06$, $t_9 = 15.09$, $p < 0.001$), and gender ($\Delta\tau = 0.01$, $t_9 = 3.17$, $p = 0.011$). By contrast, ASD-specific features correlated better with DSM IV behavioral subtypes ($\Delta\tau = 0.01$, $t_9 = 2.34$, $p = 0.044$), suggesting that CVAE identified population-wide patterns of neuroanatomy, some of which are shared by all participants and some of which are only present in those with ASD.

Additionally, the SFARI VIP dataset allowed us to ask whether neuroanatomical differences observed in 16p11.2 deletion and duplication carriers are consistent with patterns of variation in the typically developing population or whether they match patterns of variation within ASD. Similarity between deletion and duplication CNVs was better reflected in ASD-specific features than in shared features ($\Delta\tau = 0.05$, $t_9 = 14.54$, $p < 0.001$). We note that the neuroanatomical phenotypes associated with these CNVs are likely only a subset of ASD more broadly: More than 200 CNVs have been associated with autism (*1*, *5*). The future development of larger genotyped datasets will be crucial for further advances.

The nature of variation

Researchers have debated whether individual differences in ASD are better understood as distinct subtypes or as variation along continuous dimensions (*6*). Having identified ASD-specific features makes it possible to test these hypotheses directly. We used Gaussian mixture modeling to identify clusters of subjects based on each set of features, selecting the

optimal number of clusters using the Bayesian information criterion (BIC) (Fig. 1D).

Because CVAEs and VAEs are probabilistic, we determined the optimal number of clusters for 100 samples of the latent features (see SM). The subjects' VAE features were consistent with a single cluster in 100% of samples ($p < 0.01$; Fig. 1D). The CVAE results were more nuanced. For shared features, 100% of samples indicated multiple clusters ($p < 0.01$). However, the subject distribution based on the ASD-specific features again suggested continuous variation, with 100% of samples indicating a single cluster ($p < 0.01$). Thus, the results of cluster analysis show that once disentangled from typical variation, ASD-related neuroanatomical variation is better captured by continuous dimensions rather than by discrete categories. This conclusion applies to the neuroanatomical data considered here; other datasets (e.g., functional imaging datasets) might reveal multiple clusters.

Neuroanatomical interpretation

To identify loci of anatomical variation between ASD subjects, we followed a three-step process. First, for each ASD participant, we reconstructed their brain using only the "shared" features that represent individual variation that is independent of diagnosis. (Technically, we set the ASD-specific feature values to zero before using the CVAE decoder.) The result is a "synthetic TC twin": a simulated brain matched to the original ASD participant but lacking any features that our analyses identified as ASD specific. This synthetic twin is effectively a data-driven case control. In the second step, we estimated a nonrigid transformation that morphs the counterfactual TC brain to match the corresponding ASD participant's brain. This produced a vector field that described the differences between the ASD brain and the corresponding TC brain (see SM). Finally, we calculated the Jacobian determinant of the vector field. This measure captures the local volumetric compression and expansion needed to morph the simulated TC brain into the corresponding ASD brain. Repeating this procedure for all participants, we computed interpretable gray and white matter alterations that vary across the ASD participant population.

To organize the search of interpretable neuroanatomical features, we calculated the first two principal components (PCs) of the Jacobian maps across all ASD participants ($N = 470$). We then measured systematic variation in the compression and expansion of different brain regions along each PC by computing, for each voxel, the correlation between the PC loadings for that voxel and the Jacobian determinants (Fig. 2; maps thresholded at $p < 0.05$, Bonferroni corrected). By focusing on the two PCs that account for most variance (~20%), we simplify interpretation and reduce the number of comparisons. We note that although detecting intensity contrasts is comparatively more difficult in some areas (e.g., thalamus), this is a common feature for both TC and ASD brains; attendant variation should be captured by the shared features, not ASD-specific features.

To test the correspondence between the anatomical PCs and behavioral symptoms in different cognitive domains, we correlated the PC loadings with scores in ADOS communication, ADOS social, and ADOS stereotyped behaviors. The first PC positively correlated with the ADOS communication instrument ($\tau_{342} = 0.09$, $p = 0.017$) and with the stereotyped behavior instrument ($\tau_{283} = 0.10$, $p = 0.023$) but not with the ADOS social instrument ($\tau_{343} = 0.06$, $p = 0.136$). The second PC correlated positively with the ADOS communication instrument ($\tau_{342} = 0.08$, $p = 0.039$) but not with the ADOS repetitive behavior ($\tau_{283} = -0.06$, $p = 0.155$) or social instruments ($\tau_{343} = -0.04$, $p = 0.259$). A limitation of this analysis is that it relies on relatively coarse measures of behavior. Finer-grained measures of behavior that cover a broad range of cognitive abilities will be needed to identify relationships between anatomical dimensions and more-specific symptoms. This could help clarify, for instance, the importance of the volumetric changes to areas related to social cognition in the second PC (fig. S9) and of volumetric changes to Broca's area (left inferior frontal gyrus) in the first PC.

Previous work has found neuroanatomical differences between ASD participants and TCs that vary with age (*18*–*20*), and earlier in this text we reported that ASD-specific features do indeed correlate with age (Fig. 1). However, this was not the largest source of ASD-specific individual differences: The first two neuroanatomical PCs were not related to age (PC1: $\tau_{468} = 0.06$, $p = 0.064$; PC2: $\tau_{468} = 0.04$, $p = 0.159$). Clarifying age-related differences within ASD will require more-sensitive analyses, perhaps involving longitudinal data, which can have more precision for detecting age-related differences.

Discussion

These results demonstrate that disentangling ASD-specific variation in neuroanatomy from shared variation reveals correlations between individual differences at the level of brain structure and differences in symptoms as well as genetics. We find that ASD-specific features can be disentangled using a data-driven approach (CVAEs) that generalizes to new datasets without the need for additional training. This property facilitates its application in diagnostic settings, in which a model trained on previous cases can be used to analyze the data from new individuals. Note that these results represent a floor: Even more powerful models trained on larger datasets and higher-resolution inputs may identify additional, more subtle patterns. Although in this study we used CVAEs to analyze anatomical data in the context of ASD, the approach is broadly applicable to other data modalities (e.g., behavioral data, functional imaging) and to other psychiatric disorders.

Individual variation within ASD was better captured by continuous dimensions than by multiple distinct clusters, indicating that—at least at the level of neuroanatomy—dimensional approaches can provide a better account of individual variation than discrete diagnostic categories. It remains possible, however, that functional neuroimaging or genetic data will reveal clusters that are not apparent in the anatomical data.

Previous work has demonstrated that anatomical changes associated with ASD vary across different ages (*18*–*20*). Here, we found that age correlates not only with anatomical features shared with typical controls but also to some extent with ASD-specific features, consistent with the existence of ASD-specific patterns of age-dependent changes in anatomy. Multiple possible causes of volumetric changes have been hypothesized in previous studies, including differences in cell proliferation (*21*) or in soma size and dendrite length (*22*). Clarifying the structural causes and functional consequences of volumetric changes remains a critical open question in human neuroscience.

REFERENCES AND NOTES

1. D. Moreno-De-Luca, C. L. Martin, *Curr. Opin. Genet. Dev.* **68**, 71–78 (2021).
2. D. L. Christensen *et al*., *MMWR Surveill. Summ.* **65**, 1–23 (2018).
3. American Psychiatric Association, *Diagnostic and Statistical Manual of Mental Disorders (DSM-5)* (American Psychiatric Publishing, 2013).
4. S. Zheng, K. A. Hume, H. Able, S. L. Bishop, B. A. Boyd, *Autism Res.* **13**, 796–809 (2020).
5. J. Y. An, C. Claudianos, *Neurosci. Biobehav. Rev.* **68**, 442–453 (2016).
6. S.-J. Hong *et al*., *Biol. Psychiatry* **88**, 111–128 (2020).
7. S. Georgiades, P. Szatmari, M. Boyle, *Neuropsychiatry* **3**, 123–125 (2013).
8. R. Higdon *et al*., *OMICS* **19**, 197–208 (2015).
9. J. Gu, R. Kanai, *Front. Hum. Neurosci.* **8**, 262 (2014).
10. G. Auzias, S. Takerkart, C. Deruelle, *IEEE J. Biomed. Health Inform.* **20**, 810–817 (2016).
11. K. A. Severson, S. Ghosh, K. Ng, *Proc. Conf. AAAI Artif. Intell.* **33**, 4862–4869 (2019).
12. A. Abid, J. Zou, arXiv:1902.04601 [cs.LG] (2019).
13. A. Di Martino *et al*., *Mol. Psychiatry* **19**, 659–667 (2014).
14. N. Kriegeskorte, M. Mur, P. Bandettini, *Front. Syst. Neurosci.* **2**, 4 (2008).
15. S. A. Bedford *et al*., *Mol. Psychiatry* **25**, 614–628 (2020).
16. The Simons Vip Consortium, *Neuron* **73**, 1063–1067 (2012).
17. A. D'Amour *et al*., arXiv:2011.03395 [cs.LG] (2020).
18. E. Courchesne, R. Carper, N. Akshoomoff, *JAMA* **290**, 337–344 (2003).
19. T. Nickl-Jockschat *et al*., *Hum. Brain Mapp.* **33**, 1470–1489 (2012).
20. E. Greimel *et al*., *Brain Struct. Funct.* **218**, 929–942 (2013).
21. M. C. Marchetto *et al*., *Mol. Psychiatry* **22**, 820–835 (2017).
22. A. Deshpande *et al*., *Cell Rep.* **22**, 2678–2687 (2017).

23. A. Aglinskas, SCCN Lab, sccnlab/pub-CVAE-MRI-ASD: Code release for reproducibility. Zenodo (2022); https://doi.org/10.5281/zenodo.6304004.

ACKNOWLEDGMENTS

We thank M. Ritchey, L. Young, and R. Saxe for comments on a previous draft. **Funding:** This work was supported by a grant from SFARI (award no. 614379 to J.K.H. and S.A.) and by start-up funds from Boston College to S.A. **Author contributions:** Conceptualization: A.A., J.K.H., S.A.; Methodology: A.A., S.A.; Formal analysis: A.A.; Funding acquisition: J.K.H., S.A.; Supervision: S.A.; Writing – original draft: A.A., J.K.H., S.A.; Writing – review and editing: A.A., J.K.H., S.A. **Competing interests:** The authors declare no competing interests. **Data and materials availability:** The ABIDE I data is available at fcon_1000.projects.nitrc.org/indi/abide. The SFARI VIP data (now renamed Simons Searchlight) is available at base.sfari.org. All code is available at github.com/sccnlab/pub-CVAE-MRI-ASD and through Zenodo (*23*). All other data are in the main paper or supplementary materials.

SUPPLEMENTARY MATERIALS

science.org/doi/10.1126/science.abm2461
Materials and Methods
Figs. S1 to S9
Tables S1 to S3
References (*24*–*39*)
MDAR Reproducibility Checklist

View/request a protocol for this paper from *Bio-protocol*.

Submitted 29 September 2021; accepted 26 April 2022
10.1126/science.abm2461

Science

Contrastive machine learning reveals the structure of neuroanatomical variation within autism

Aidas AglinskasJoshua K. HartshorneStefano Anzellotti

Science, 376 (6597), • DOI: 10.1126/science.abm2461

Brain structure in ASD

Autism spectrum disorder (ASD) may be characterized by impaired social interactions, but persons with ASD also struggle with a variety of other behavioral and intellectual difficulties. Are individual differences better understood as ASD subtypes or as continuous variation? Aglinskas *et al*. analyzed magnetic resonance imaging brain scans to look for brain differences that can be attributed to ASD and not to other causes of individual variation. The authors found evidence for continuous variation and identified two axes of variation in brain structure. Such clarity about ASD variation may help to fine-tune interventions for individual patients. —PJH