NeuroImage 226 (2021) 117606

Contents lists available at ScienceDirect

NeuroImage

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

Emergence and organization of adult brain function throughout child development

Tristan S. Yates*, Cameron T. Ellis, Nicholas B. Turk-Browne

Department of Psychology, Yale University, New Haven, CT 06520, USA

ARTICLE INFO

Keywords:
Brain maturation
Child development
fMRI
Machine learning
Naturalistic perception
Shared response modeling
Developmental neuroscience

ABSTRACT

Adult cognitive neuroscience has guided the study of human brain development by identifying regions associated with cognitive functions at maturity. The activity, connectivity, and structure of a region can be compared across ages to characterize the developmental trajectory of the corresponding function. However, developmental differences may reflect both the maturation of the function and also its organization across the brain. That is, a function may be present in children but supported by different brain regions, leading its maturity to be underestimated. Here we test the presence, maturity, and localization of adult functions in children using shared response modeling, a machine learning approach for functional alignment. After learning a lower-dimensional feature space from fMRI activity as adults watched a movie, we translated these shared features into the anatomical brain space of children 3–12 years old. To evaluate functional maturity, we correlated this reconstructed activity with children's actual fMRI activity as they watched the same movie. We found reliable correlations throughout cortex, even in the youngest children. The strength of the correlation in the precuneus, inferior frontal gyrus, and lateral occipital cortex predicted chronological age. These age-related changes were driven by three types of developmental trajectories: emergence from absence to presence, consistency in anatomical expression, and reorganization from one anatomical region to another. We also found evidence that the processing of pain-related events in the movie underwent reorganization across childhood. This data-driven, naturalistic approach provides a new perspective on the development of functional neuroanatomy throughout childhood.

1. Introduction

The advent of non-invasive neuroimaging techniques opened a new window into the study of human cognitive development. Initial fMRI studies of children examined functional differences in anatomical brain regions associated with particular cognitive functions in adults, such as the prefrontal cortex for executive function (Luna et al., 2001) and the amygdala for fear processing (Thomas et al., 2001). This approach was effective in characterizing the development of these brain regions. It also provided evidence in support of a maturational account of development (Johnson, 2011), which states that as cognitive functions come online during development, they will occupy the same neural regions as adults.

However, other research has shown that cognitive functions can be subserved by different brain regions at different ages (Bayet and Nelson, 2019; Brown et al., 2005; Durston et al., 2006; Jolles et al., 2011; Nelson et al., 2003; Schlaggar, 2002; Thomason et al., 2008). One striking example is the development of visual object recognition: face processing is initially supported by both left and right fusiform gyrus, but as children learn to read, a region selective for visual words emerges in the left fusiform gyrus and face processing begins to shift to being right lateralized (Centanni et al., 2018; Dehaene-Lambertz et al., 2018; Dundas et al., 2013). Thus, as children gain new cognitive skills, such as reading, the localization of orthographic (and face) processing changes.

There are multiple interpretations for dynamic patterns of cognitive development in the brain (Brown et al., 2006; Poldrack, 2010). For example, studies sometimes show a shift from distributed to focal processing over development (Durston et al., 2006); that is, functions are localized to many regions early on, but localized to one or a smaller number of regions later. Such a finding could be evidence of increased efficiency of brain regions, decreased reliance on other regions for "support," changes in the computations being performed, or simply artifacts of greater variability in region localization in developing populations.

Changes in the number and relationship between brain regions supporting a cognitive function can also be considered in the context of the interactive specialization framework (Johnson, 2001; 2011). This theory emphasizes that brain regions do not mature in isolation, but specialize over experience through interactions with each other. For example, if a newly emerging skill or function would be well-supported by a brain region (e.g., because of its cells, circuitry, or connectivity) that currently supports a different function, and if there is another brain region that already supports or could support the current function, there

* Corresponding author.
E-mail address: tristan.yates@yale.edu (T.S. Yates).

https://doi.org/10.1016/j.neuroimage.2020.117606
Received 6 July 2020; Received in revised form 21 November 2020; Accepted 25 November 2020
Available online 30 November 2020

may be changes in tuning and connectivity between these regions to accommodate the new function while retaining the current function in conjunction with other regions.

The current study builds on prior work exploring dynamic changes in functional brain development in two ways. Hypothesis-driven approaches can find evidence for the maturational account, but have a harder time supporting alternatives. An analysis that can discover the reorganization of a function on the cortex would be valuable in discovering how the brain matures functionally. We therefore establish an unbiased, data-driven approach that can capture different kinds of developmental change under a common framework, by searching for patterns of functional activity rather than focusing on specific regions. We do not assume the timing or types of cognitive functions that are recruited (though consider this in a secondary analysis), but instead rely on the idea that the time course of brain activity reflects the computations being performed. This means that similar time courses of brain activity can be attributed to similar cognitive processes. Second, most brain-based studies of cognitive development pursue a standard lab approach of isolating cognitive functions. Although vital for manipulating and tracking exactly what is changing, an alternative, naturalistic approach could provide a more comprehensive and ecological sense of how the brain is developing. Indeed, in adult cognitive neuroscience, naturalistic paradigms such as movie-watching have yielded unexpected insights into how the brain processes information across time-scales and domains (Sonkusare et al., 2019). Movie-watching has also emerged as an invaluable tool in children (see Vanderwal et al., 2019), with previous studies tending to focus on one or a small number of key cognitive functions (Cantlon and Li, 2013; Richardson et al., 2018). Here we operationalize function more holistically, as a collection of data-driven features derived from adult brain activity. In this way, movie-watching can be used to efficiently sample a broad swath of cognition.

To track the neural development of cognitive functions within and across brain regions, we applied functional alignment (Chen et al., 2015) to an open-access dataset ("Partly Cloudy") of children aged 3–12 and adults watching a movie during fMRI (Richardson et al., 2018). We used open-source software for shared response modeling (SRM; Kumar et al., 2020) to extract temporal features of brain activity that were shared across the adults. For meaningful features to be extracted, this method requires that brain activity is time-locked across participants, as occurs when they watch the same movie (Hasson et al., 2004). Therefore, SRM is currently not well-suited for resting-state fMRI, for which there is no expectation that spontaneous activity will be aligned in time. There are larger movie datasets from children than Partly Cloudy (Alexander et al., 2017), which present additional opportunities for future research. However, to our knowledge, the Partly Cloudy dataset is unique in having both children across a range of ages and an adult comparison group who watched the same movie. This is crucial for the present goal of learning features of adult function and assessing their expression in children as a way of quantifying development.

In prior work, SRM has been used to distinguish scenes during movie-watching (Chen et al., 2015; Turek et al., 2018), relate perception and recall (Chen et al., 2017), and map semantic features to fMRI activity (Vodrahalli et al., 2018). Here, we take a new approach of using SRM to understand how content in the adult brain is represented in the developing brain. Although a variant of SRM has been used for age prediction amongst adults ranging in age from 18 to 88 (Richard et al., 2019), our study is unique in learning a shared response from one age group (adults) and applying it to a completely independent age group (children). By studying a large sample across childhood, this approach can be used to characterize the developmental trajectory of adult brain function. The features that SRM learns can be thought of as capturing abstract cognitive functions that vary distinctively from each other across the movie in a way that is consistent in adults. We then mapped the children into this lower-dimensional feature space. These mappings were used in reverse to port adult fMRI activity into each child's anatomical brain space. Comparing this reconstructed activity to the child's actual fMRI activity allowed us to quantify the expression of adult functions throughout childhood. Higher correlation between reconstructed and actual activity means that the child's brain expressed the abstract functions shared in adults, what we refer to as "adult-like". Although these abstract functions cannot be cleanly identified with specific psychological constructs, a key advantage of this data-driven approach is that they can be aligned across adults and children without making any anatomical assumptions. There is no requirement that functions are instantiated in the same brain regions across individuals, whether within or between ages. In fact, our approach would be equally sensitive to the development of functions that emerge within one region as to functions that reorganize from one region to another.

2. Materials and methods

2.1. Data

The Partly Cloudy dataset was obtained from the OpenNeuro database (accession number ds000228). A full description of data acquisition can be found in the original paper (Richardson et al., 2018). Participants with neuroimaging data available consisted of 33 adults (18–39 years old; $M = 24.8$, $SD = 5.3$; 20 female) and 122 children (3.5–12 years old; $M = 6.7$, $SD = 2.3$; 64 female; for more details, see Table S1). Informed consent was obtained from adult participants and from parents/guardians on behalf of child participants, who provided their own assent. The study was approved by the Committee on the Use of Humans as Experimental Subjects (COUHES) at the Massachusetts Institute of Technology.

2.2. fMRI acquisition and preprocessing

Participants watched an animated movie (Sohn and Reher, 2009) that lasted approximately 5 minutes while undergoing fMRI. No explicit task was given beyond staying still and paying attention to the movie. Adults and older children used the standard Siemens 32-channel head coil. For younger children, one of two custom 32-channel phased-array head coils was used (smallest coil: $N = 3$, $M = 3.91$, $SD = 0.42$ years old; smaller coil: $N = 28$, $M = 4.07$, $SD = 0.42$ years old). The only difference between head coils was their size. These size-optimized head coils have been shown to increase signal-to-noise in participants with smaller heads (Keil et al., 2011). fMRI data were collected using a gradient-echo EPI sequence (TR = 2 s, TE = 30 ms, flip angle = 90°, matrix = 64×64, slices = 32, interleaved slice acquisition) covering the whole brain. To correct for slight variations in the voxel size and slice gap parameters across participants, data were resampled to 3 mm isotropic with 10% slice gap (the modal parameters). Children also participated in a number of behavioral tasks not related to the movie that are beyond the scope of the current study.

Preprocessing of the structural and functional MRI data was performed with fMRIPrep (v1.1.8; Esteban et al., 2019). First, T1-weighted structural images from an MPRAGE sequence (GRAPPA = 3, slices = 176, resolution = 1 mm isotropic, adult coil FOV = 256 mm, child coils FOV = 192 mm) were corrected for intensity non-uniformity using N4BiasFieldCorrection (v2.1.0) and skull-stripped using antsBrainExtraction.sh (v2.1.0, OASIS template). Cerebrospinal fluid (CSF), white-matter (WM) and gray-matter (GM) masks were extracted from the structural image using FAST (FSL v5.0.9). Surface reconstruction was performed by FreeSurfer (v6.0.1). Nonlinear registration to an MNI template for spatial normalization was performed with the antsRegistration tool (ANTs v2.1.0). Registrations were visually inspected and the quality of fit did not seem to differ across child and adult participants. Functional images were slice-time corrected using 3dTshift from AFNI (v16.2.07), then motion corrected using FSL's mcflirt (v5.0.9). Co-registration to the structural scan was performed with 9 degrees of freedom using bbregister in FreeSurfer (v6.0.1). Transformations were con-

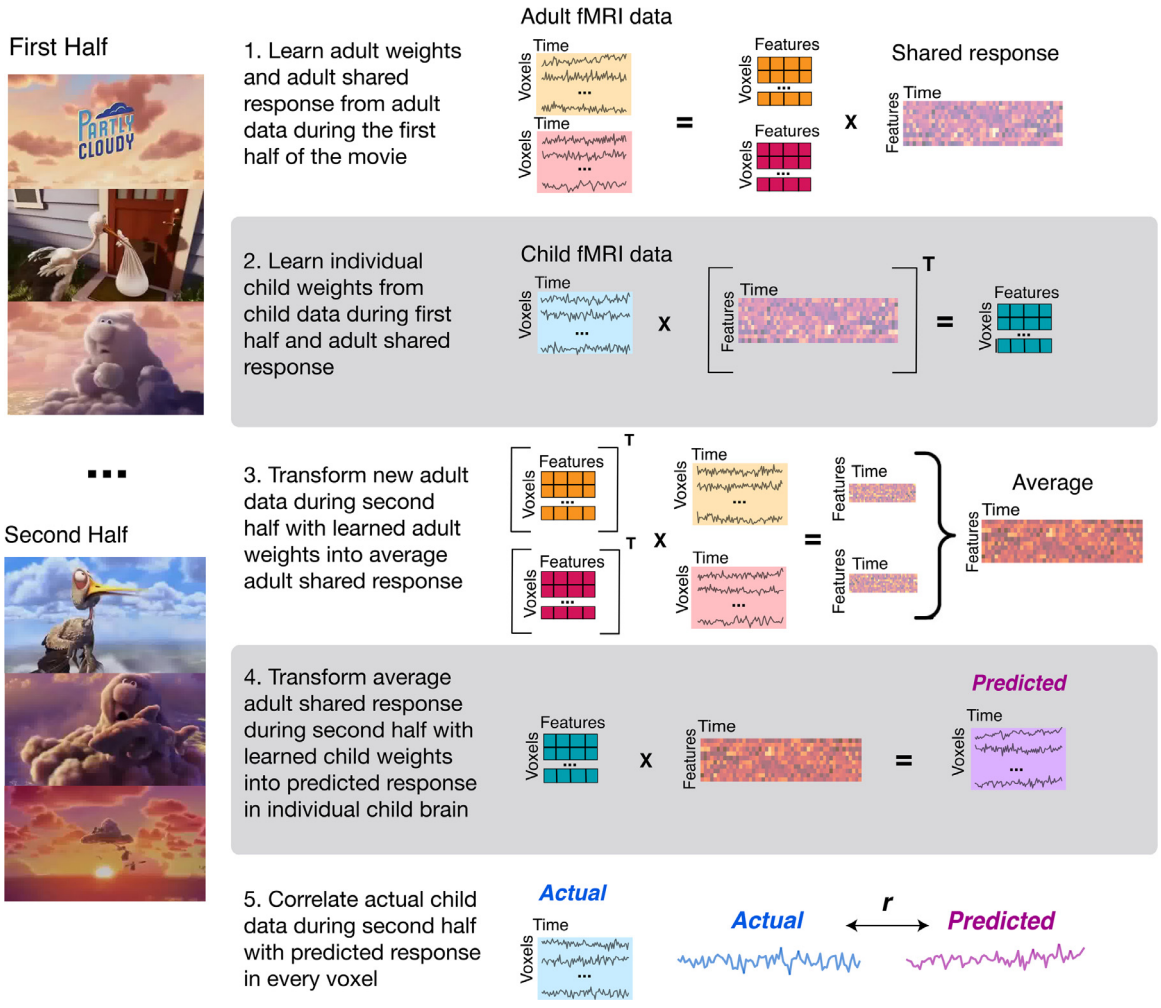

Fig. 1. Schematic of the signal reconstruction pipeline. An SRM was trained on the first half of the fMRI data from a group of adults (N = 33; 2 example matrices shown) and then each of the children (N = 122) was fit into this space. Adult fMRI data from the second half of the movie (i.e., not used to train the model) were transformed into the shared space and averaged. This shared adult activity was then projected into each child's brain and correlated with their actual activity. This procedure was then repeated for training on the second half and testing on the first half.

catenated using antsApplyTransform (v2.1.0). Frame-wise displacement was estimated using Nipype.

2.3. Experimental design

We used SRM (Chen et al., 2015; Turek et al., 2018) to identify activity in the developing brain that could be predicted from adult brain activity (illustrated in Fig. 1). This method assumes that all participants were shown the same stimulus with the same number of time-points but does not require that they have the same number of voxels. First, the time-points from a group of adults were evenly split into training and test sets for cross-validation purposes. We used one half of the adult data to learn the shared feature space, consisting of features that captured shared temporal variance across adults, as well as the mappings between individual adults' brain activity and this shared space. No child data were used for training the model. Prior to any other analyses, we ran this analysis on subsets of the adult data varying the number of features (5–80, in increments of 5) and found that 10 features learned from a set of adults gave the highest whole-brain signal reconstruction values for held-out adults (M = 0.087, SD = 0.031; Fig. S1). Although this was the global maximum, other numbers of features yielded comparable signal reconstruction. Selecting one of these local maxima would change the dimensionality of the shared response, which could affect the results. This would be unwieldy to examine in the current paper, but could be explored more thoroughly in future work, including by sampling different numbers of features with more granularity in steps of 1 rather than 5 features.

After learning 10 shared features in adults using one half of the adult data, we found the mapping (voxels by features) between an individual child's functional activity (voxels by time) and the shared response (time by features) for this same portion of the movie. Singular value decomposition was implemented to solve for the orthogonal weight matrix. Values in each cell of this resulting weight matrix denote how strongly a given voxel in the child expresses each of the 10 features discovered from the adult data. Next, we used the remaining half of the adult data to quantify how the 10 shared features were expressed in data not involved in SRM training. Each adult's transposed weight matrix (features by voxels) was used to transform their raw voxel activity (voxels by time) into the shared feature space (features by time). We then averaged these shared responses across all adults to find the canonical adult response in terms of shared features during this part of the movie. Finally, each child's weight matrix (voxels by features) was used to transform the average adult shared response (features by time) into the child's brain space (voxels by time). This predicted response represents what the child's brain activity would look like if they expressed the same shared features of adults. We quantified the extent to which this was true by correlating the child's actual raw response with this predicted response for each voxel separately. Thus, higher signal reconstruction

reflected greater adult-child functional similarity—i.e., more adult-like functions in the child's brain—agnostic to the anatomical localization of these functions in either group. The voxelwise map of predicted-actual correlations for each child was averaged across individuals within age groups. We ran this entire procedure twice, training the SRM on the first half of the movie and testing on the second half in one fold, and then vice versa in another fold, and present results averaged across these folds. Rather than using Pearson correlation as a standardized measure of similarity, we could have estimated this association with a general linear model (GLM) for each voxel, in which the predicted activity from adults is entered as the explanatory variable and the actual child activity is treated as the response variable. The resulting beta values could be used to test reliability across participants, which yields nearly identical statistical significance as correlation coefficients (Fig. S2), but with the added benefit of reflecting the real magnitude of the relationship between these variables. The GLM approach also makes it possible to control for nuisance variables, such as head motion.

Our first objective was to assess the degree to which children's brain activity could be reconstructed from shared features learned in adults. We quantified the noise ceiling for this group-level signal reconstruction by leaving one adult participant out of SRM training, correlating that individual's predicted and actual brain activity, and then iterating through each adult. This was treated as the noise ceiling because the held-out participant was from the same age group used to train the SRM. Our second objective was to quantify how signal reconstruction may change over development, and whether this could be a useful measure for predicting an entirely held-out child's age. We then explored how the individual features that comprise the shared response may exhibit different developmental trajectories throughout childhood. Finally, we investigated the relationship between individual features and cognitive constructs. Partly Cloudy was first used as a localizer for theory of mind in adults (Jacoby et al., 2016). As such, different events in the movie relevant to social cognition were annotated, including social, pain, and mentalizing events. For each event type, we generated a time series of events and convolved it with a double-gamma hemodynamic response function (HRF). Because we used data from the second half of the movie for predicting children's brain activity from adult features, we restricted our analyses of social cognitive events in that half. Two of the three event types, pain and mentalizing events, were present in the second half of the movie and could therefore be compared to the average expression of shared features in adults.

2.4. Statistical analysis

We used bootstrap resampling methods to statistically evaluate our results non-parametrically (Efron and Tibshirani, 1986; Fan et al., 2020; Kim et al., 2014). For each effect of interest, at the last step of the analysis we randomly sampled participants with replacement to form a new sample of the same size as the original group, averaged the effect across the sample, and repeated for 10,000 iterations. The logic of this approach is that if an effect is reliable across participants, the participants should be interchangeable, and a similar group effect should be observed in each iteration. The resampled values across all iterations reflect a sampling distribution of the effect of interest, further providing confidence intervals on the original effect. Null hypothesis testing can be performed by determining the proportion of resampled values that were of the opposite sign as the original effect. The original effect can also be normalized into a *z*-statistic by dividing the mean of the resampled distribution by its standard deviation. For voxelwise analyses, this was performed in each voxel to create a statistical map. This map was corrected for multiple comparisons using a cluster-based correction in FSL's cluster tool (cluster-forming threshold, $p<0.001$). Corrected *p*-values were found using Gaussian Random Field Theory and the smoothness estimated from the original map.

We quantified the relationship between signal reconstruction and age by first fitting a linear regression model for each voxel. We then used the same bootstrapping approach described above, now resampling participants to calculate the relationship between signal reconstruction and age in each iteration, resulting in a sampling distribution for the relationship. We calculated the *p*-value as the proportion of iterations on which the correlation coefficient from the linear regression model went in the opposite direction from the original model. We compared this model against other types of models that have been used previously in developmental cognitive neuroscience (Schlichting et al., 2017). For each voxel, we fit five regression models: (1) a linear model with age alone as the predictor (as above), (2) a linear model with age and sex as predictors, (3) a linear model with age and sex as predictors plus an age-by-sex interaction term, (4) a quadratic model with just age as a predictor, and (5) a quadratic model with age and sex as predictors plus an age-by-sex interaction term. We then assessed which model gave the lowest Akaike information criterion (AIC), a measure of the relative quality of different models, for each voxel.

We used leave-one-out cross-validation to predict the age of children from signal reconstruction. For each iteration, we fit the linear regression model between signal reconstruction and age in a training set of *N*-1 participants. We did this separately for each voxel and retained the clusters that were significant within the training set (based on the previously described bootstrap resampling method). Note that we ignored the sign of the significant relationship in a cluster and thus it was possible to find negative beta values. We then fit a regularized ridge regression (penalty = 1) across voxels from the significant clusters. To predict age in the held-out *N*th test participant, we input their signal reconstruction scores across these voxels and output an estimated age. Finally, we calculated the Pearson correlation and mean-squared error between the chronological and predicted ages of children across iterations.

In addition to reconstructing all 10 adult features in children, we also performed signal reconstruction for individual adult features. To test single-feature reconstruction within the adults, we could not perform the fully cross-validated approach described above of leaving one adult participant entirely out of both the training set used to learn the SRM and the testing set used to generate predicted activity. This is because each training set would have contained a unique set of adults, which could lead to different features and/or a different ordering of features in the shared space. We would therefore not be confident that we were considering the same feature across folds. Instead, we included all adults when training the SRM on one half of the movie, so that there would be a consistent shared space across adults and as used to reconstruct children. Nevertheless, we left one adult out when averaging the adult shared response for the other half of the movie, using the expression of the selected feature in all but that adult to predict their neural activity. Because the reconstructed adult was used in SRM training, we included a 10 time-point buffer between their training and test data to minimize non-independence. Signal reconstruction of individual features in children was identical to the main analysis, except based on separate weight matrices for each child mapping from their voxel space to a given adult feature. That is, although the adult data could not be fully cross-validated, the data from children remained completely untouched during SRM training.

In the single-feature analysis, we sought to quantify how brain regions changed in their expression of features over development. We thus defined regions of interest (ROIs) using the Schaefer brain atlas (Schaefer et al., 2018). This atlas consists of 100 parcels discovered from resting-state connectivity data in adults and matched to 17 functional networks (Yeo et al., 2011). We reconstructed each of the ten features in adults and children and then calculated the average signal reconstruction scores across voxels in each of the 100 parcels. For statistical analysis, we used the same bootstrap resampling procedure across the participants in a given age group, separately for each parcel and feature. To correct for multiple comparisons across the parcels, we used Bonferroni correction (100,000 bootstrap iterations were run to gain precision on *p*-values for thresholding). Finally, the parcels that survived correction were ranked according to the strength of signal reconstruction.

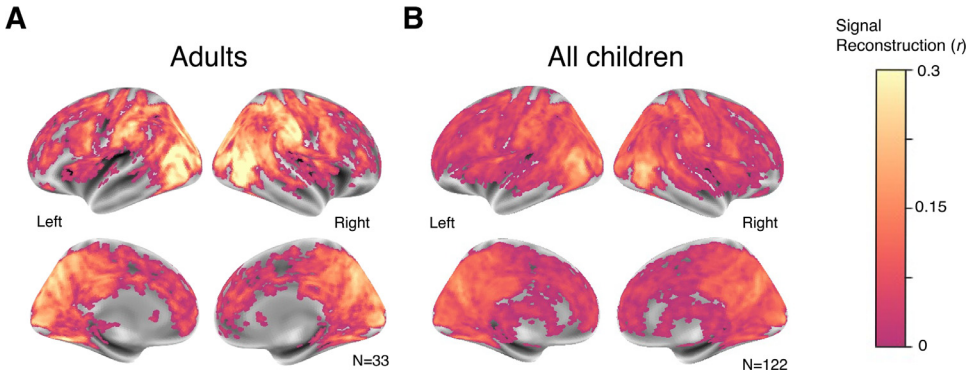

Fig. 2. Reconstruction of adult function in children. In all brain plots, the strength of signal reconstruction is denoted by color and only regions that survived statistical thresholding through cluster correction are plotted. (A) Signal reconstruction for a group of adults predicting an independent adult's functional activity is reliable throughout much of the brain. (B) Signal reconstruction remains statistically reliable (though numerically weaker) for adults predicting functional activity in children ranging from 3–12 years old.

We used bootstrap resampling to quantify the relationship between individual features and psychological constructs. For each feature, we randomly sampled with replacement the expression of the feature in 33 adults (i.e., second half data transformed into their feature space), and then averaged the results into a single adult shared response for the feature. We then correlated the average adult shared response for each feature with the convolved event time series for each event type and repeated the procedure 10,000 times. We calculated the *p*-value as the proportion of resamples in which the correlation had the opposite sign from the original correlation, doubled to convert to two-tailed. Because this was not a planned analysis, we corrected for 10 multiple comparisons (corresponding to the 10 features) with a Bonferroni correction.

2.5. Code accessibility

The analysis code for running the signal reconstruction analysis pipeline is available on Github: https://github.com/tristansyates/partly-recon.

3. Results

3.1. Adult-like brain function in early to middle childhood

We first characterized how well adult brain activity could be reconstructed from other adults. Signal reconstruction was widespread throughout the brain, especially in occipital and parietal cortices (Fig. 2A). This indicates that the shared features learned by SRM from one half of the movie accounted for adult brain activity during the other half of the movie. Importantly, this only works because of the ability of SRM to learn abstract features that generalize across the contents of the two halves. Although we found tentative evidence for higher whole-brain signal reconstruction in younger adults (Fig. S3), scores were reliably positive across the adult sample.

Remarkably, signal reconstruction was also widespread in the children despite the fact that the functional features were defined entirely in adults (Fig. 2B). Results were nearly identical when including nuisance parameters in a GLM relating children's actual and predicted brain activity (Fig. S2) and were similar, but slightly stronger after accounting for nuisance parameters prior to constructing the shared response (Fig. S4). In the child brain, adult functions were most strongly represented in lateral occipital and posterior medial regions, albeit weaker than in the adult brain.

3.2. Relationship between age and signal reconstruction

The previous analysis collapsed across all children, but the degree and location of signal reconstruction may vary with age. We quantified these relationships by correlating, for each voxel in the brain, children's signal reconstruction values (i.e., the extent to which adult brain activity predicts that child's brain activity) with their chronological age. After correcting for multiple comparisons, signal reconstruction was positively correlated with children's age in regions including the bilateral precuneus, bilateral lateral occipital cortex, postcentral gyrus, and inferior frontal gyrus (Fig. 3A; see also Fig. S3). Thus, in these regions, adult brain activity better predicts older children's brain activity compared to younger children. No regions showed a reliable negative correlation. Alternative models taking into account children's sex and testing for quadratic relationships did not generally provide better fits than this linear model (Table 1). The basic linear model with age alone gave the lowest AIC values for the majority of voxels, and therefore minimized the information loss when trading off with model complexity. Furthermore, individual voxels in which other models had the lowest AIC values were scattered across the brain, suggesting that they were capturing noise and providing further evidence that the basic linear model performed best.

Table 1
Different models used to predict signal reconstruction (SigRecon), which is the correlation between predicted and actual activity in a held-out child. For each voxel, the model with the lowest AIC was assigned to that voxel. The average and standard deviation of AIC values for a given model for voxels where that model was the best is shown in the right column. Overall, the linear model used in our main analyses with age as the only predictor best described the data.

Regression model	Number of voxels	AIC *M* (*SD*)
$SigRecon \sim Age$	39,111	2.41 (0.56)
$SigRecon \sim Age + Sex$	60	5.21 (1.39)
$SigRecon \sim Age + Sex + Age * Sex$	31	4.76 (1.91)
$SigRecon \sim Age^2$	49	3.75 (2.01)
$SigRecon \sim Age^2 + Sex^2 + Age * Sex$	31	3.71 (1.05)

3.3. Out-of-sample prediction of a child's age from signal reconstruction

With chronological age related to signal reconstruction in several regions of the brain, it may also be possible to predict the age of a previously unseen child. In a nested cross-validation analysis, we first trained a linear regression model between signal reconstruction and age for each voxel in all but one child. Blind to this child, we determined which voxels showed a significant relationship with age again through bootstrapping and cluster correction. We then trained a ridge regression model on these significant voxels. This model was used to predict the held-out child's chronological age from their multivariate pattern of signal reconstruction scores across the voxels. This procedure was repeated 122 times to use each child as the held-out test data once. Note that the significant clusters varied slightly across iterations because the training set changed when different children were used as test data. Finally, we correlated the predicted and actual ages (Fig. 3B), and found a strong relationship ($r = 0.436$, $p < 0.001$). Indeed, our model had a mean-squared error of 6.05, meaning that our average error in age prediction was 2.46 years across an age range of 8.78 years.

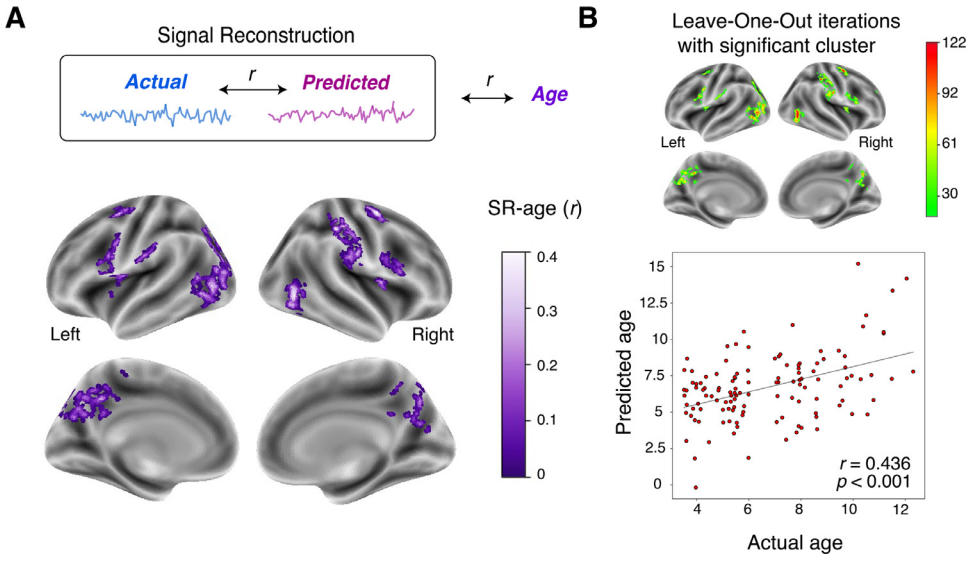

Fig. 3. Relationship between signal reconstruction and age. (A) Brain regions with a reliable correlation between signal reconstruction of adult function in a child's brain activity and the child's chronological age across all 122 children, colored by the strength of the relationship. (B) Similar regions are found in leave-one-child-out iterations of the age prediction analysis. Yellow–red colors signify regions that were significant in a majority of iterations. Using signal reconstruction scores from these regions, we could accurately predict the held-out child's chronological age.

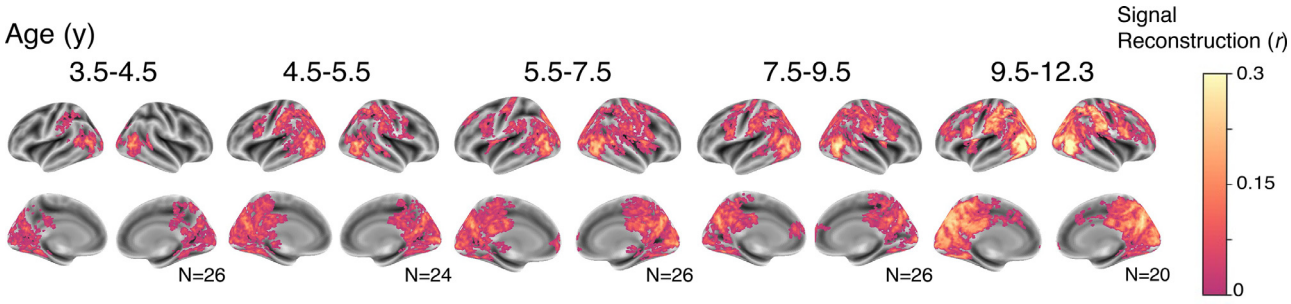

Fig. 4. Signal reconstruction of adult features was statistically reliable even in the youngest children, but spread anatomically and grew in strength throughout childhood. To quantify this developmental change, we correlated the unthresholded voxelwise signal reconstruction in each age group with that of adults, revealing increasing maturity: 3.5–4.5 years, $r = 0.507$; 4.5–5.5 years, $r = 0.587$; 5.5–7.5 years, $r = 0.601$; 7.5–9.5 years, $r = 0.646$; 9.5–12.3, $r = 0.797$.

3.4. Reliable signal reconstruction in all age groups

The relationship between signal reconstruction and age could reflect a lack of adult function in early childhood that emerges in middle childhood. To evaluate this possibility, we divided children into five age bins (3.5–4.5, 4.5–5.5, 5.5–7.5, 7.5–9.5, 9.5–12.3 years old), each containing roughly the same number of participants (N = 20–26). This was done for analytical convenience and was not intended to suggest discrete developmental stages. Although signal reconstruction increased with age, we nevertheless found reliable signal reconstruction in every age group. This includes lateral occipital, posterior medial, and supramarginal regions, even in the youngest children aged 3–5 (Fig. 4). Signal reconstruction emerged in frontal regions around age 5, and became more pronounced in the older groups. To obtain a global measure of adult-child similarity, we correlated the unthresholded maps of signal reconstruction for each age group with that of adults. There was reasonable agreement in all groups, though the amount of variance explained grew from 25% in the youngest children to 64% in the oldest children.

3.5. Controlling for age-related noise in signal reconstruction

Increases in signal reconstruction over development may result from younger children being "noisier" than older children and adults, including because of differences in task compliance, preprocessing quality, and/or BOLD physiology (Harris et al., 2011; Phan et al., 2018). Children did move their heads more than adults overall, but this did not track with age across children (analysis of number of time-points exceeding 2 mm motion threshold from Richardson et al. (2018): one-way ANOVA across age groups, $F(4,116) = 1.175$, $p = 0.325$; correlation with age across children, $r = -0.112$, $p = 0.221$).

Moreover, we can estimate and control for noise in different age bands using the noise-ceiling approach from adults (Fig. 2A). For each age group, we held one child out and used SRM to learn shared features in the remaining children of that group. We then predicted the held-out child's voxel activity, correlated it with their actual activity, and averaged across significant clusters to derive a global within-group signal reconstruction score for each child. The average score across children in a group provides a measure of the reliability of functional brain activity in that group. This within-group signal reconstruction correlated with chronological age ($r = 0.359$, $p < 0.001$), consistent with decreasing noise over development. Within-group signal reconstruction was also correlated with adult-group signal reconstruction ($r = 0.647$, $p < 0.001$). Critically, however, the correlation between adult-group signal reconstruction and chronological age ($r = 0.418$, $p < 0.001$) persisted after controlling for within-group signal reconstruction ($r = 0.261$, $p = 0.003$). In contrast, the correlation between within-group signal reconstruction and chronological age did not hold after accounting for adult-group signal reconstruction ($r = 0.128$, $p = 0.158$). These results suggest that adult features capture more about child brain function than changes in noise over development.

3.6. Emergence and reorganization of adult function over child development

There are at least two other potential explanations for the age-related increases in signal reconstruction we observed. First, a subset of the

adult functions being reconstructed may be absent from younger children and mature over development to become present in older children (emergence). Second, adult functions may be present in both younger and older children but expressed in different brain regions over development (reorganization). By both accounts, the brain regions expressing a function in older children would not express it as strongly in younger children. The accounts differ, however, in that reorganization but not emergence predicts that the function would be expressed in other brain regions in younger children.

Emergence and reorganization are difficult to distinguish with the analysis approach used so far. By predicting the activity of each voxel as a weighted combination of all adult features, we may have obscured developmental trajectories that differed across features. We thus modified our pipeline to predict activity in children from individual adult features, each of which captures a narrower, more unique range of adult function. The modification occurred in step 4 (Fig. 1), where we now transformed only one average adult feature at a time into the voxel space of a child. Individual features do not necessarily isolate single functions, and emergence and reorganization are not mutually exclusive, so it may be possible to observe both patterns within a feature. We used a functional atlas (Schaefer et al., 2018) to identify regions that showed the strongest signal reconstruction for a given feature per age group.

We found evidence of both emergence and reorganization across different adult features, as well as a third pattern in which a feature was expressed in the same brain region(s) across development (consistency). Representative features illustrating these three types of trajectories are depicted in Fig. 5 (for all features, see Figs. S5 and S6). For example, Feature 4 was not reliably expressed in the two youngest age groups and emerged in the lingual gyrus of older children and adults (Fig. 5A). In contrast, Feature 6 was expressed most strongly in the posterior cingulate and lingual gyrus consistently throughout development (Fig. 5B). Finally, Feature 7 was expressed most strongly in the precuneus and posterior cingulate of children and migrated to be more strongly represented in parietal regions in adults (Fig. 5C). This feature also interestingly shows some consistency over development. Nonetheless, it was one of several features where the average signal reconstruction value across the whole brain was significantly related to child age (Fig. S3). Thus, by measuring functional profiles regardless of anatomy, signal reconstruction revealed developmental changes both within and across brain regions.

3.7. Cognitive interpretation of shared features

The features that we learned from adult fMRI activity are abstractions, making it difficult to assign them to specific cognitive functions. Moreover, our data-driven approach with SRM means that it is possible that some cognitive functions may be partially distributed across features, while others may not account for enough variance to be included in the model. Nonetheless, in a follow-up analysis, we explored the relationship of these features to cognitive functions that were evoked by this movie. An earlier study annotated different events relevant to social cognition in the same movie, including pain and mentalizing event types (Jacoby et al., 2016). For each event type, we convolved the time series of events with a double-gamma HRF and correlated it with the average expression of each shared feature from adults (Fig. 6). For pain events, only one of the features (Feature 9) was reliably correlated ($p < 0.05$, corrected). For mentalizing events, none of the features were correlated. Interestingly, Feature 9 was most strongly expressed in the cuneus of adults and the oldest children, the postcentral gyrus near the temporo-parietal junction of the middle age group (eight- and nine-year-olds) and the posterior cingulate of younger children (Fig. S6). These findings in adults and children over eight are in line with the prior studies showing that the cuneus is more active for pain than mentalizing events in adults (Jacoby et al., 2016) and that bilateral postcentral gyrus is a node in the pain network (Bruneau et al., 2015; 2012). Although the posterior cingulate is usually more activated for mentalizing events than pain events in adults, here it is related to pain processing in young children. These results highlight that the localization of pain processing in the developing brain is dynamic, with the role of the posterior cingulate and other regions changing during this time. Furthermore, this shows that even quite young children are capable of representing the pain state of others, and that applying data-driven then confirmatory analyses can be a powerful combination for understanding cognitive development in the brain.

4. Discussion

In this study, we sought to bring a new perspective to the long-standing question of how and when the developing brain becomes "adult-like" (Johnson, 2011; Somerville, 2016). The typical approach for answering this question is to align children and adults into a common anatomical space and compare activity between groups in the same brain regions (Cantlon and Li, 2013; Dosenbach et al., 2010; Fair et al., 2009; Gogtay et al., 2004; Richardson et al., 2018). Thus, even when the goal is to understand functional similarities and differences over development, anatomy serves as a guide and constraint. The alternative approach we employed is to align children and adults into a common *functional* space, which allowed us to quantify adult-like brain activity in children without making any assumptions about a consistent mapping between function and anatomy over development. This anatomically agnostic approach has the advantage of finding representations in the developing brain that may otherwise be overlooked. It does not require pre-specifying the type of function that is expected to differ across development *a priori*. Instead, it uses a data-driven approach to extract content from the movie that explains brain activity and to identify where in the brain this content is represented. In children as young as 3.5 years old watching a short movie, we found regions of the brain, especially in occipital cortex, that reliably expressed functional features shared amongst adults who watched the same movie. Based on where and how strongly these features were expressed, we were able to build a predictive model of age that depended only on brain activity during movie watching. We then demonstrated the power of functional alignment by revealing features of adult function that emerge and reorganize across anatomical locations over development. Finally, we showed how confirmatory hypothesis testing can be performed within this framework to interpret shared adult functions and how they develop.

We interpreted increasing signal reconstruction with age as evidence of functional specialization and maturation in the developing brain. A related but slightly different framing is that brain functionality itself was not always changing in these cases, but rather it was the way that children deployed this functionality during the movie. For instance, if older children attended to the content of the movie in a more adult-like fashion than younger children, this may have affected perceptual input to downstream functions and increased similarity to adult brain activity. The defining characteristic of this interpretation is that younger children may possess the *capacity* for such functions but not engage them because of attentional differences in perceptual input. Even if attention was allocated similarly across age, richer schematic knowledge in older children may have enhanced their understanding of the movie narrative (Brod et al., 2017; Ghosh and Gilboa, 2014) by highlighting connections between objects and events that may not otherwise be easily integrated. Again, younger children may have the capacity for this kind of integration in principle but be unable to deploy it without access to the relevant conceptual knowledge. Of course, attention and memory are functions of the brain, and so developmental differences in these functions are what we sought to characterize in the first place. The key point is that increasing signal reconstruction could reflect the development of a function or of necessary perceptual or conceptual precursors to that function.

Signal reconstruction allowed us to build a predictive model of age based on how strongly children's brains represented the shared features of adults. Our predictive model was highly significant, but a limitation

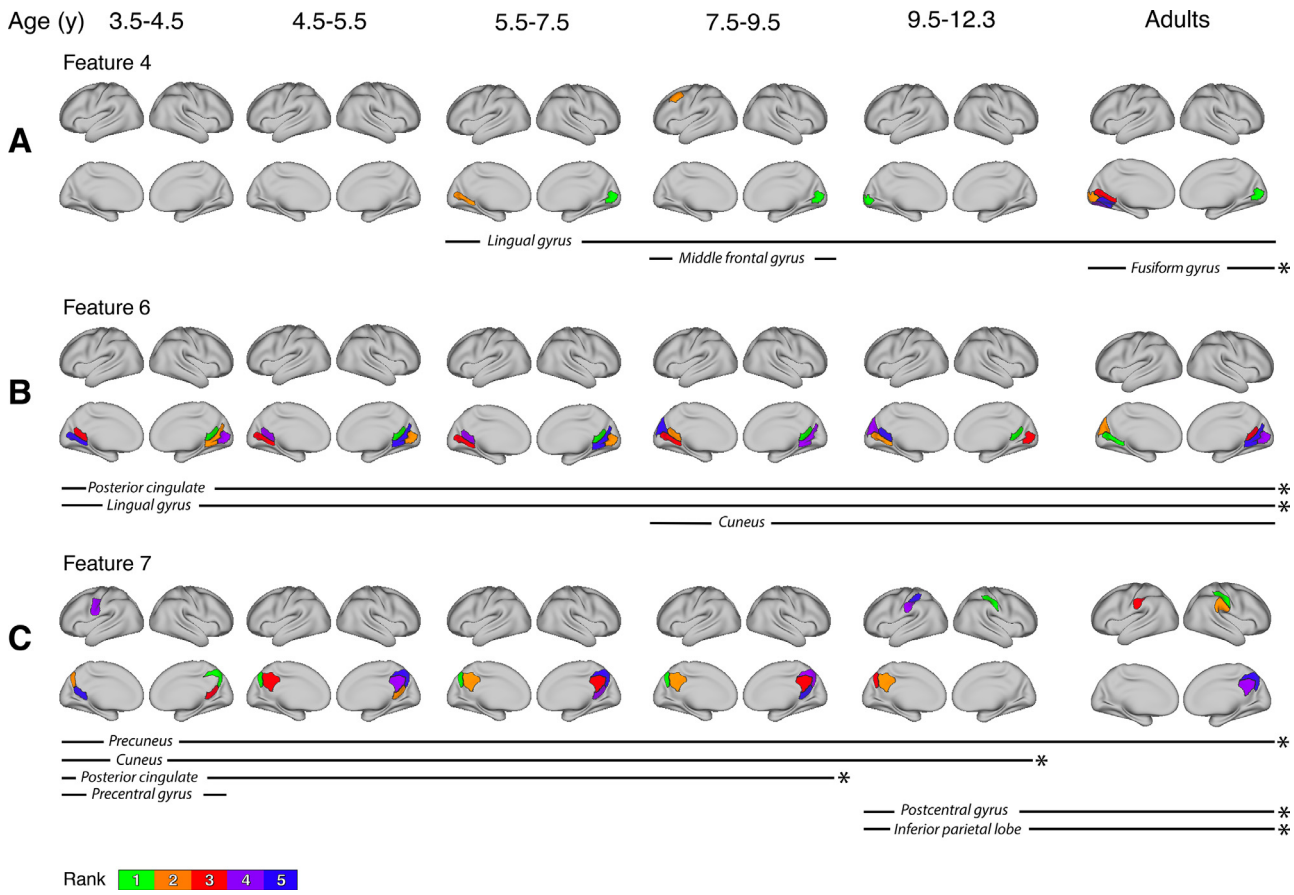

Fig. 5. Trajectories of functional development within and across brain regions. (A–C) To understand the nature of developmental changes in signal reconstruction, we predicted activity from one adult feature at a time rather than all features. We used a functional parcellation to identify which regions expressed a given feature most strongly in each age group. Parcels with significant signal reconstruction of adult features within each age group ($p<0.05$, corrected) were ranked by the strength of the reconstruction. For ease of visualization, here we color up to the top five parcels for each feature and age group. The anatomical labels for these parcels were obtained from the Talairach atlas. Three example adult features are depicted across ages, illustrating developmental trajectories we refer to as emergence (Feature 4), consistency (Feature 6), and reorganization (Feature 7). The top five parcels for the remaining features are depicted in Fig. S5 and all parcels that are significant for each group are displayed in Fig. S6. Asterisks indicate which parcels differed significantly in signal reconstruction of each feature across age groups (p <0.05, corrected; Fig. S7).

is that the strength of the relationship was moderate, and the mean error in age was substantial for the sample's age range. Larger sample sizes and more training data (i.e., longer movie) in future studies could increase precision. Nonetheless, other studies using out-of-sample cross-validation methods like ours have found a similar range of relationship strength (Finn et al., 2015; Lin et al., 2018).

A variant of the signal reconstruction approach allowed us to identify different types of developmental trajectories across adult features. We used ROIs defined by functional connectivity (Schaefer et al., 2018) to map the developmental trajectories of neural features shared amongst adults. The voxels that comprise a region in this atlas have homogeneous functional activity and connectivity in adults. We used the term "emergence" to describe features of adult brain activity that were not reliably expressed in any regions of young children's brains, but appeared in older children and were present in the same location up until adulthood. Features that showed "consistency" were those in which the localization remained consistent in all of the age groups tested, including the youngest children. Finally, features that showed "reorganization" were those that were reliably expressed in at least one parcel in the brains of children and adults, but where the localization of these parcels varied across ages. Therefore, features that exhibit the first trajectory (emergence) may comprise late-developing cognitive functions, while the other feature types comprise cognitive functions that can be represented by younger children.

Although both emergence and consistency of features in adult regions over development are consistent with multiple accounts of brain development, the reorganization of features over development cannot be explained by a maturational account, which argues that certain cognitive functions are tied to particular brain regions and minimally influenced by the environment and nearby regions. This may apply to certain highly specialized regions, such as for vision or language (Kanwisher, 2010), but our results highlight that many adult features are not characterized by a one-to-one mapping between structure and function, and that assuming this might obscure functional similarities across development. It is worth noting that a more pure form of reorganization, whereby a feature in adulthood is no longer expressed at all in the regions in which it was previously expressed, was less common in our study. Instead, we tended to observe *relative* reorganization, whereby the set of regions expressing a feature remains fairly consistent over development but the rank order of which regions show the strongest expression changes. For example, the precuneus was replaced by the postcentral gyrus as the region with highest signal reconstruction of feature 7 for children around 9 years old, even though the precuneus continued to express the function through adulthood. Future work tracking individuals longitudinally should try to understand why reorganization occurs and how it relates to environmental changes or new skill acquisition (e.g., reading; Dehaene-Lambertz et al., 2018).

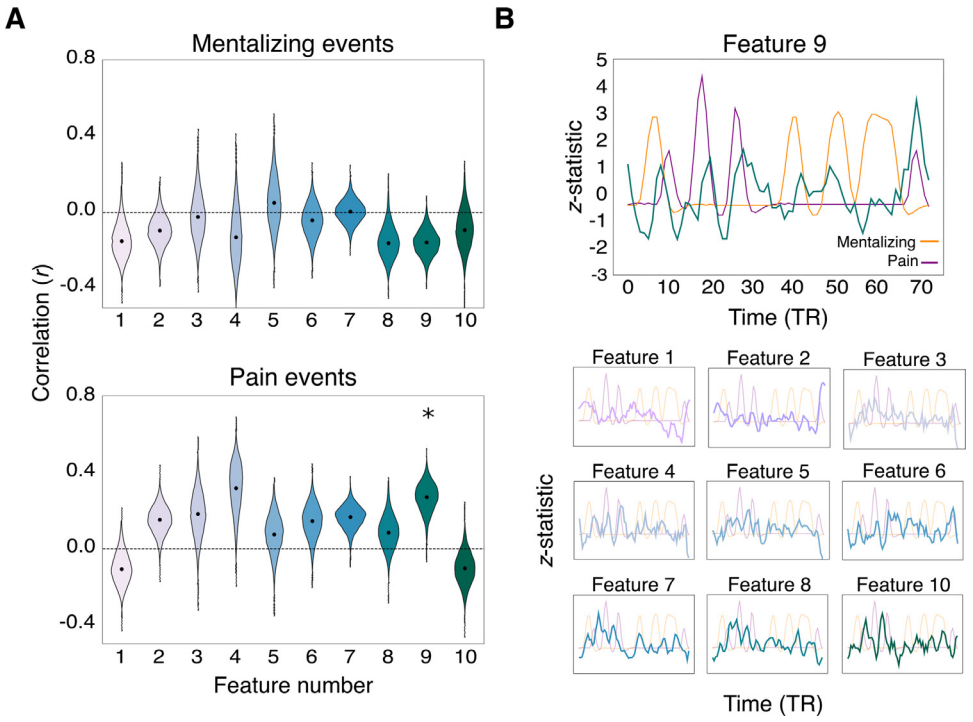

Fig. 6. Distribution of resampled correlation values between the average representation of shared features in a group of adults and two of the three cognitive constructs defined in Jacoby et al. (2016). (A) For mental events, we did not find a significant relationship with any of the shared features of adults. However, we did find a significant positive correlation between Feature 9 and pain events (p <0.05, corrected). This feature was strongly expressed in the cuneus of adults and older children, the postcentral gyrus in the middle child age group, and the posterior cingulate of younger children. (B) Visualization of the z-scored time courses of individual features during the second half of the movie and the two cognitive constructs. Orange lines represent the HRF-convolved mentalizing events, and purple lines represent the HRF-convolved pain events. The feature time course lines are colored similarly to (A).

Our findings that several relevant features are present as early as 3 years of age, in either the same or different regions as expressed in adults, suggest early adultlike cognition. Even so, these data-driven features remain abstract and are not easily decomposed into specific cognitive functions. Moreover, although the features captured unique and substantial variance shared across adults, each may still embed multiple cognitive functions with similar temporal profiles of brain activity. This has implications for interpreting features showing anatomical reorganization of function over development (e.g., Fig. 5C). Specifically, our definition of reorganization was that the same functions were subserved by different brain regions over development—that is, a cognitive function that manifests in region X of younger children is expressed in region Y of adults. This could occur if the original region was co-opted by a different function (Behrmann and Plaut, 2015) or if the nature of the function changed with increasing skill and expertise (Johnson, 2001).

However, the possibility of multiple functions being embedded in a given feature suggests an alternative interpretation. Namely, these functions may have stable organization over development, but the relative weighting of the functions as captured by the feature may change. Consider a hypothetical feature that is active during the title and credits of the movie. This feature might capture multiple language-related functions engaged by these scenes, such as letter recognition in region X and semantic comprehension in region Y. We would expect even the younger age groups to respond to the orthography of the words and thus show signal reconstruction in X, but perhaps only the older children and adults would respond to the meaning of the words and show signal reconstruction in Y. Disentangling these possibilities requires a better understanding of how the abstract features from SRM relate to the contents of the movie and to the cognitive functions that are engaged. Future studies could make progress in this direction by using reverse correlation (Hasson et al., 2004) or hand-coded events in the movie to better ascertain the functional profile of the features. Indeed, we found that Feature 9 was related to the processing of pain events identified by prior annotations of the movie. This pain-related feature demonstrates the power of our data-driven approach to understanding cognitive development, as it was expressed in different regions of the brain in younger children and adults. Although the regions recruited in older children and adults were predicted by previous research (Bruneau et al., 2015; Jacoby et al., 2016), the region recruited in younger children, the posterior cingulate, is typically associated with mentalizing rather than observing physical pain (Bruneau et al., 2015; Saxe and Powell, 2006). Thus, signal reconstruction allowed us to find evidence for commonalities in the ability to process pain events over development, despite differences in anatomical localization. Although annotations were not available for other types of events, this movie likely engaged other functions related to visual processing, object recognition, and narrative comprehension. The relationship between adult shared features and these cognitive functions remains an avenue for future research.

Future work could also address the cognitive underpinnings of shared features by selecting or designing movies to target specific cognitive functions. Indeed, a constraint in our study, and SRM more generally, is that the features extracted depend on the movie. The use of other data for functional alignment, including from live action videos, different sensory modalities, or synchronized trials of varied cognitive tasks could sample cognition even more broadly. This might allow SRM to learn a richer functional space that provides a more complete picture of functional brain development across childhood. Additionally, it might reveal other types of developmental trajectories that were not evident in the current study. In one of the three types of neural trajectories we defined (reorganization), a function that may be present behaviorally from a young age undergoes neural changes, such that it is subserved by one region early on before reorganizing to another region later in development. The transition between these two regions may inform behavioral findings of a U-shaped (or inverted U-shaped) curve, where younger children and adults are more similar than children of intermediate ages (Siegler, 2004). Combining this approach with behavioral measures could therefore reveal why such changes occur.

Another limitation of the current study is that we rely on shared features learned in adults, yet there may be developmental changes within the adult cohort. Indeed, the adult sample includes a large age range, and we found some evidence that younger adults had higher signal reconstruction than older adults (Fig. S3). However, the younger adults comprise a larger proportion of the sample, which likely biased the SRM features to be more consistent with their features. Regardless, any adult heterogeneity does not compromise our analyses of children; signal reconstruction within the adults was reliable across the sample, suggesting

that the shared response was able to capture functions that are stable over this age range. One takeaway from this result is that our approach can be applied successfully to adults (akin to Richard et al., 2019). It also hints at the possibility of learning shared features in younger age groups and testing on older age groups. Thus, in addition to characterizing the emergence of adult features of cognition over development, our analyses could be applied in reverse to answer the complementary question of what child features *disappear* over development into adulthood (or in aging from young adults to the elderly). Future work with larger cohorts of similarly-aged children may be able to answer this question.

We focused on brain development, but the techniques in our paper could be applied productively to a number of questions that involve comparing functional activity across groups. For instance, learning the functional features shared amongst a clinical population and then reconstructing these features in an undiagnosed individual may be useful for predicting whether the individual will develop the condition. This method could also be used to assess how and when a learner's brain starts to resemble that of an expert over the course of training. Because signal reconstruction does not require that the group and individual have the same brain sizes or even anatomical organization, this approach could even be applied between humans and non-human animals to trace how cognitive functions are shared over phylogeny. Indeed, there is no requirement that the group and individual be brains at all, which could, for example, allow states of a computational model to be ported into the brain for model-based analysis, or vice versa for brain-computer interfaces.

Declaration of Competing Interest

The authors declare no competing financial interests.

Credit authorship contribution statement

Tristan S. Yates: Conceptualization, Methodology, Software, Formal analysis, Writing - original draft, Writing - review & editing. **Cameron T. Ellis:** Conceptualization, Methodology, Software, Writing - original draft, Writing - review & editing. **Nicholas B. Turk-Browne:** Conceptualization, Methodology, Writing - original draft, Writing - review & editing, Supervision, Funding acquisition.

Acknowledgments

T.S.Y was supported by NSF Graduate Research Fellowship DGE 1752134. N.B.T-B. was supported by NIH R01 MH069456, NSF CCF 1839308, and the Canadian Institute for Advanced Research. We are grateful for the effort and generosity of Hilary Richardson and the rest of the research team that originally collected this dataset in Rebecca Saxe's lab at MIT. We also thank lab members for helpful feedback on earlier versions of the project.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at 10.1016/j.neuroimage.2020.117606.

References

Alexander, L.M., Escalera, J., Ai, L., Andreotti, C., Febre, K., Mangone, A., Vega-Potler, N., Langer, N., Alexander, A., Kovacs, M., Litke, S., O'Hagan, B., Andersen, J., Bronstein, B., Bui, A., Bushey, M., Butler, H., Castagna, V., Camacho, N., Chan, E., Citera, D., Clucas, J., Cohen, S., Dufek, S., Eaves, M., Fradera, B., Gardner, J., Grant-Villegas, N., Green, G., Gregory, C., Hart, E., Harris, S., Horton, M., Kahn, D., Kabotyanski, K., Karmel, B., Kelly, S.P., Kleinman, K., Koo, B., Kramer, E., Lennon, E., Lord, C., Mantello, G., Margolis, A., Merikangas, K.R., Milham, J., Minniti, G., Neuhaus, R., Levine, A., Osman, Y., Parra, L.C., Pugh, K.R., Racanello, A., Restrepo, A., Saltzman, T., Septimus, B., Tobe, R., Waltz, R., Williams, A., Yeo, A., Castellanos, F.X., Klein, A., Paus, T., Leventhal, B.L., Craddock, R.C., Koplewicz, H.S., Milham, M.P., 2017. An open resource for transdiagnostic research in pediatric mental health and learning disorders. Sci. Data 4 (1), 170181.

Bayet, L., Nelson, C.A., 2019. The perception of facial emotion in typical and atypical development. In: LoBue, V., Pérez-Edgar, K., Buss, K.A. (Eds.), Handbook of Emotional Development. Springer International Publishing, pp. 105–138.

Behrmann, M., Plaut, D.C., 2015. A vision of graded hemispheric specialization. Ann. N. Y. Acad. Sci. 1359, 30–46.

Brod, G., Lindenberger, U., Shing, Y.L., 2017. Neural activation patterns during retrieval of schema-related memories: differences and commonalities between children and adults. Dev. Sci. 20 (6), e12475.

Brown, T.T., Lugar, H.M., Coalson, R.S., Miezin, F.M., Petersen, S.E., Schlaggar, B.L., 2005. Developmental changes in human cerebral functional organization for word generation. Cereb. Cortex 15 (3), 275–290.

Brown, T.T., Petersen, S.E., Schlaggar, B.L., 2006. Does human functional brain organization shift from diffuse to focal with development? Dev. Sci. 9 (1), 9–11.

Bruneau, E.G., Jacoby, N., Saxe, R., 2015. Empathic control through coordinated interaction of amygdala, theory of mind and extended pain matrix brain regions. NeuroImage 114, 105–119.

Bruneau, E.G., Pluta, A., Saxe, R., 2012. Distinct roles of the 'shared pain' and 'theory of mind' networks in processing others' emotional suffering. Neuropsychologia 50 (2), 219–231.

Cantlon, J.F., Li, R., 2013. Neural activity during natural viewing of sesame street statistically predicts test scores in early childhood. PLOS Biol. 11 (1), e1001462.

Centanni, T.M., Norton, E.S., Park, A., Beach, S.D., Halverson, K., Ozernov. Palchik, O., Gaab, N., Gabrieli, J.D., 2018. Early development of letter specialization in left fusiform is associated with better word reading and smaller fusiform face area. Dev. Sci. 21 (5), e12658.

Chen, J., Leong, Y.C., Honey, C.J., Yong, C.H., Norman, K.A., Hasson, U., 2017. Shared memories reveal shared structure in neural activity across individuals. Nat. Neurosci. 20 (1), 115–125.

Chen, P.-H.C., Chen, J., Yeshurun, Y., Hasson, U., Haxby, J., Ramadge, P.J., 2015. A reduced-dimension fMRI shared response model. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (Eds.), Advances in Neural Information Processing Systems 28. Curran Associates, Inc., pp. 460–468.

Dehaene-Lambertz, G., Monzalvo, K., Dehaene, S., 2018. The emergence of the visual word form: longitudinal evolution of category-specific ventral visual areas during reading acquisition. PLOS Biol. 16 (3), e2004103.

Dosenbach, N.U.F., Nardos, B., Cohen, A.L., Fair, D.A., Power, J.D., Church, J.A., Nelson, S.M., Wig, G.S., Vogel, A.C., Lessov-Schlaggar, C.N., Barnes, K.A., Dubis, J.W., Feczko, E., Coalson, R.S., Pruett, J.R., Barch, D.M., Petersen, S.E., Schlaggar, B.L., 2010. Prediction of individual brain maturity using fMRI. Science 329 (5997), 1358–1361.

Dundas, E.M., Plaut, D.C., Behrmann, M., 2013. The joint development of hemispheric lateralization for words and faces. J. Exp. Psychol. Gen. 142 (2), 348–358.

Durston, S., Davidson, M.C., Tottenham, N., Galvan, A., Spicer, J., Fossella, J.A., Casey, B.J., 2006. A shift from diffuse to focal cortical activity with development. Dev. Sci. 9 (1), 1–8.

Efron, B., Tibshirani, R., 1986. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. Stat. Sci. 1 (1), 54–75.

Esteban, O., Markiewicz, C.J., Blair, R.W., Moodie, C.A., Isik, A.I., Erramuzpe, A., Kent, J.D., Goncalves, M., DuPre, E., Snyder, M., Oya, H., Ghosh, S.S., Wright, J., Durnez, J., Poldrack, R.A., Gorgolewski, K.J., 2019. fMRIPrep: a robust preprocessing pipeline for functional MRI. Nat. Methods 16 (1), 111–116.

Fair, D.A., Cohen, A.L., Power, J.D., Dosenbach, N.U.F., Church, J.A., Miezin, F.M., Schlaggar, B.L., Petersen, S.E., 2009. Functional brain networks develop from a "local to distributed" organization. PLoS Comput. Biol. 5 (5), e1000381.

Fan, J.E., Wammes, J.D., Gunn, J.B., Yamins, D.L.K., Norman, K.A., Turk-Browne, N.B., 2020. Relating visual production and recognition of objects in human visual cortex. J. Neurosci. 40 (8), 1710–1721.

Finn, E.S., Shen, X., Scheinost, D., Rosenberg, M.D., Huang, J., Chun, M.M., Papademetris, X., Constable, R.T., 2015. Functional connectome fingerprinting: Identifying individuals based on patterns of brain connectivity. Nat. Neurosci. 18 (11), 1664–1671.

Ghosh, V.E., Gilboa, A., 2014. What is a memory schema? A historical perspective on current neuroscience literature. Neuropsychologia 53, 104–114.

Gogtay, N., Giedd, J.N., Lusk, L., Hayashi, K.M., Greenstein, D., Vaituzis, A.C., Nugent, T.F., Herman, D.H., Clasen, L.S., Toga, A.W., Rapoport, J.L., Thompson, P.M., 2004. Dynamic mapping of human cortical development during childhood through early adulthood. Proc. Natl. Acad. Sci. 101 (21), 8174–8179.

Harris, J.J., Reynell, C., Attwell, D., 2011. The physiology of developmental changes in BOLD functional imaging signals. Dev. Cognit. Neurosci. 1 (3), 199–216.

Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., Malach, R., 2004. Intersubject synchronization of cortical activity during natural vision. Science 303 (5664), 1634–1640.

Jacoby, N., Bruneau, E., Koster-Hale, J., Saxe, R., 2016. Localizing pain matrix and theory of mind networks with both verbal and non-verbal stimuli. NeuroImage 126, 39–48.

Johnson, M.H., 2001. Functional brain development in humans. Nat. Rev. Neurosci. 2 (7), 475–483.

Johnson, M.H., 2011. Interactive specialization: a domain-general framework for human functional brain development? Dev. Cognit. Neurosci. 1 (1), 7–21.

Jolles, D.D., Kleibeuker, S.W., Rombouts, S.A.R.B., Crone, E.A., 2011. Developmental differences in prefrontal activation during working memory maintenance and manipulation for different memory loads. Dev. Sci. 14 (4), 713–724.

Kanwisher, N., 2010. Functional specificity in the human brain: a window into the functional architecture of the mind. Proc. Natl. Acad. Sci. 107 (25), 11163–11170.

Keil, B., Alagappan, V., Mareyam, A., McNab, J.A., Fujimoto, K., Tountcheva, V., Triantafyllou, C., Dilks, D.D., Kanwisher, N., Lin, W., Grant, P.E., Wald, L.L., 2011. Size-optimized 32-channel brain arrays for 3 t pediatric imaging: Pediatric brain arrays. Magn. Reson. Med. 66 (6), 1777–1787.

Kim, G., Lewis-Peacock, J.A., Norman, K.A., Turk-Browne, N.B., 2014. Pruning of memories by context-based prediction error. Proc. Natl. Acad. Sci. 111 (24), 8997–9002.

Kumar, M., Ellis, C.T., Lu, Q., Zhang, H., Capotă, M., Willke, T.L., Ramadge, P.J., Turk-Browne, N.B., Norman, K.A., 2020. BrainIAK tutorials: user-friendly learning materials for advanced fMRI analysis. PLOS Comput. Biol. 16 (1), e1007549.

Lin, Q., Rosenberg, M.D., Yoo, K., Hsu, T.W., O'Connell, T.P., Chun, M.M., 2018. Resting-state functional connectivity predicts cognitive impairment related to Alzheimer's disease. Front. Aging Neurosci. 10, 94. https://pubmed.ncbi.nlm.nih.gov/29706883/.

Luna, B., Thulborn, K.R., Munoz, D.P., Merriam, E.P., Garver, K.E., Minshew, N.J., Keshavan, M.S., Genovese, C.R., Eddy, W.F., Sweeney, J.A., 2001. Maturation of widely distributed brain function subserves cognitive development. NeuroImage 13 (5), 786–793.

Nelson, E.E., McClure, E.B., Monk, C.S., Zarahn, E., Leibenluft, E., Pine, D.S., Ernst, M., 2003. Developmental differences in neuronal engagement during implicit encoding of emotional faces: an event-related fMRI study. J. Child Psychol. Psychiatry 44 (7), 1015–1024.

Phan, T.V., Smeets, D., Talcott, J.B., Vandermosten, M., 2018. Processing of structural neuroimaging data in young children: bridging the gap between current practice and state-of-the-art methods. Dev. Cognit. Neurosci. 33, 206–223.

Poldrack, R.A., 2010. Interpreting developmental changes in neuroimaging signals. Hum. Brain Mapp. 31 (6), 872–878.

Richard, H., Martin, L., Pinho, A. L., Pillow, J., Thirion, B., 2019. Fast shared response model for fMRI data. ArXiv:https://arxiv.org/abs/1909.12537.

Richardson, H., Lisandrelli, G., Riobueno-Naylor, A., Saxe, R., 2018. Development of the social brain from age three to twelve years. Nat. Commun. 9 (1), 1–12. https://www.nature.com/articles/s41467-018-03399-2.

Saxe, R., Powell, L.J., 2006. It's the thought that counts: specific brain regions for one component of theory of mind. Psychol. Sci. 17 (8), 692–699.

Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.-N., Holmes, A.J., Eickhoff, S.B., Yeo, B.T.T., 2018. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. Cereb. Cortex 28 (9), 3095–3114.

Schlaggar, B.L., 2002. Functional neuroanatomical differences between adults and school-age children in the processing of single words. Science 296 (5572), 1476–1479.

Schlichting, M.L., Guarino, K.F., Schapiro, A.C., Turk-Browne, N.B., Preston, A.R., 2017. Hippocampal structure predicts statistical learning and associative inference abilities during development. J. Cognit. Neurosci. 29 (1), 37–51.

Siegler, R.S., 2004. U-shaped interest in u-shaped development–and what it means. J. Cognit. Dev. 5 (1), 1–10.

Sohn, P., Reher, K., 2009. Partly Cloudy[Motion Picture].

Somerville, L.H., 2016. Searching for signatures of brain maturity: what are we searching for? Neuron 92 (6), 1164–1167.

Sonkusare, S., Breakspear, M., Guo, C., 2019. Naturalistic stimuli in neuroscience: critically acclaimed. Trends Cognit. Sci. 23 (8), 699–714.

Thomas, K.M., Drevets, W.C., Whalen, P.J., Eccard, C.H., Dahl, R.E., Ryan, N.D., Casey, B.J., 2001. Amygdala response to facial expressions in children and adults. Biol. Psychiatry 49 (4), 309–316.

Thomason, M.E., Race, E., Burrows, B., Whitfield-Gabrieli, S., Glover, G.H., Gabrieli, J.D.E., 2008. Development of spatial and verbal working memory capacity in the human brain. J. Cognit. Neurosci. 21 (2), 316–332.

Turek, J.S., Ellis, C.T., Skalaban, L.J., Turk-Browne, N.B., Willke, T.L., 2018. Capturing shared and individual information in fMRI data. In: Proceedings of the 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 826–830. ISSN: 2379-190X.

Vanderwal, T., Eilbott, J., Castellanos, F.X., 2019. Movies in the magnet: naturalistic paradigms in developmental functional neuroimaging. Dev. Cognit. Neurosci. 36, 100600.

Vodrahalli, K., Chen, P.-H., Liang, Y., Baldassano, C., Chen, J., Yong, E., Honey, C., Hasson, U., Ramadge, P., Norman, K.A., Arora, S., 2018. Mapping between fMRI responses to movies and their natural language annotations. NeuroImage 180, 223–231.

Yeo, B.T.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., Fischl, B., Liu, H., Buckner, R.L., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J. Neurophysiol. 106 (3), 1125–1165.