Characterizing the Network Architecture of Emotion Regulation Neurodevelopment

João F. Guassi Moreira[1]*, Katie A. McLaughlin[2], & Jennifer A. Silvers[1]

[1]Department of Psychology, University of California, Los Angeles, 90095
[2]Department of Psychology, Harvard University, 02138

*Corresponding author
jguassimoreira@ucla.edu

CONFLICTS OF INTEREST: The authors declare no competing financial interests

ACKNOWLEDGMENTS: We thank the members of the Harvard Stress & Development lab for their help in data collection. This research was supported by Grant MH103291 from the National Institute of Mental Health award to KAM. Preparation of the article was supported by a National Science Foundation Graduate Research Fellowship (NSF-GRF) to JFGM (Fellow ID: 2016220797).

Abstract

The ability to regulate emotions is key to goal attainment and wellbeing. Although much has been discovered about how the human brain develops to support the acquisition of emotion regulation, very little of this work has leveraged information encoded in whole-brain networks. Here we employed a network neuroscience framework in conjunction with machine learning to: (i) parse the neural underpinnings of emotion regulation skill acquisition while accounting for age, and (ii) build a working taxonomy of brain network activity supporting emotion regulation in a sample of youth (N = 70, 34 female). We were able to predict emotion regulation ability, but not age, using network activity metrics from whole-brain networks during an emotion regulation task. Further, by leveraging analytic techniques traditionally used in evolutionary biology (e.g., cophenetic correlations), we were able to demonstrate that brain networks evince reliable taxonomic organization to meet emotion regulation demands in youth. This work shows that meaningful information about emotion regulation development is encoded in whole-brain network activity, suggesting that brain activity during emotion regulation encodes unique information about regulatory skill acquisition in youth but not domain-general maturation.

Significance Statement

The acquisition of emotion regulation is critical for healthy functioning in later adult life. To date, little is known about how brain networks support the developmental acquisition of emotion regulation skills. This is noteworthy because brain networks have been increasingly shown to provide highly useful information about neural activity. Here we show that brain activity during an emotion regulation task encodes information about regulatory abilities over and above age. These results suggest emotion regulation skills are dependent on neural specialization of domain-specific systems, whereas age is encoded via domain-general systems.

Introduction

       Emotions promote adaptive behavior but can become detrimental if left unchecked (Gross, 2015), underscoring the importance of emotion regulation. Emotion regulation acquisition is a protracted process that emerges across development, and is highly consequential for wellbeing (Calkins, 1994; Casey et al., 2008; Thompson et al., 2008). Little is known about how brain networks, as opposed to individual regions of interest, support the acquisition of emotion regulation. This omission is noteworthy given adult work showing many cognitive functions are supported by brain networks (Gratton et al., 2018; Kleckner et al., 2017). Leveraging network neuroscience may help to fill two major gaps in our knowledge of emotion regulation neurodevelopment. First, this approach may improve our understanding of how correlated, but distinct, developmental phenomena (e.g., chronological age, emotion regulation ability) are represented in the brain. Second, they may reveal the underlying organization (taxonomy) of brain networks involved in emotion regulation neurodevelopment.

       Network neuroscience is a promising approach for parsing distinct but related developmental phenomena (Richardson et al., 2018). The current study focused on two network properties that have been particularly important for understanding development in other domains and thus may be fruitful for pinpointing the neurodevelopment of emotion regulation ability independent of age effects. The first property is network differentiation—the extent to which activity within a network is similar or distinct across task states, perhaps reflecting a network's "specialization" for a given task. Some networks (e.g., default mode) might support domain-general development (i.e., chronological age), whereas others might specifically support emotion regulation abilities (e.g., salience network). The second property is network community structure—or, how brain networks are divided into sub-structures ("communities", "modules"; Sporns, 2010; Yeo et al., 2011). In the context of emotion regulation, a network's modules may support different emotion regulation sub-processes (e.g., working memory). We assessed network community structure by examining the number of communities in a network as well as modularity (strength of intra- versus inter-module connections), relative to an emotional baseline condition. A network's community structure (community number, modularity) might track differentially with age and ability, helping reveal the manner in which domain-general maturation (chronological age) differs from expertise in a skill (ability). Because other research has successfully leveraged network differentiation and community structure to build models of cognitive skills and brain development (Baum et al., 2017; Richardson et al., 2018), this approach may benefit the study of emotion regulation neurodevelopment.

       Network neuroscience may also aid in another longstanding scientific goal—taxonomical organization (classification based on shared characteristics) (Linnæus, 1758). While human neuroscience has recently begun to adopt taxonomic approaches (e.g., categorizing task and resting network activity; Bolt et al., 2017; Eisenberg et al., 2018; Power et al., 2011; Yeo et al., 2011), no prior work has generated a taxonomy of brain networks involved in emotion regulation neurodevelopment. Understanding how networks developmentally organize may reveal how subcomponents of emotion regulatory processes are related in ways behavioral and univariate analyses cannot. For instance, if a network involved in emotion generation (e.g., salience) is

taxonomically similar to one involved in cognitive control (e.g., control), that would imply regulatory processes are executed by changing the stimulus' 'raw' representation. Alternately, if the two networks are taxonomically distinct, that implies regulatory cognitions act on emotional substrates via top-down processes. While taxonomies can be constructed at any developmental stage, they are important to establish in youth because taxonomies at this point likely capture information about inter-network ontogenetic relationships. A taxonomy of emotion regulation in a pediatric sample obviates speculation about ontogeny, yielding a useful baseline for adult work.

The current exploratory study had two aims. First, we used three metrics of network activity (differentiation, community number, modularity) to predict regulation ability independent of age in a sample of youth who completed emotion regulation during functional magnetic resonance imaging (fMRI). Second, we leveraged evolutionary biology techniques (agglomerative clustering, cophenetic correlation) to build a working taxonomy of networks as they pertain to emotion regulation neurodevelopment.

Methods

Participants. Participants for the current study were drawn from a larger sample of individuals enrolled in a longitudinal neuroimaging study about childhood maltreatment and affective neurodevelopment. Youth in the current analyses were selected from a non-maltreated community control group who were screened with the following exclusion criteria: childhood maltreatment/violence, presence of a developmental disorder (e.g., autism), psychotropic medication use, and IQ < 75. To qualify for inclusion in the current report, participants were required to have anatomical images free of abnormalities or artifacts, low levels of motion during the scan (addressed below), and accompanying behavioral data from our scanning emotion regulation paradigm. Our final sample included 70 youth (34 female) ranging in age from 8.08 to 17.00 years (Mean age = 12.7). The study was conducted in a large, metropolitan area in the Pacific Northwest region of the United States. Youth and their parents provided written assent and consent, respectively, in accordance with the University of Washington's Institutional Review Board. Participants were remunerated $75 after completion of the scan. Data and analysis code can be accessed on the Open Science Framework (osf.io/bfg3d).

Experimental Design

fMRI Emotion Regulation Paradigm. Participants completed a common and widely used emotion regulation task while undergoing fMRI (Silvers, Shu, Hubbard, Weber, & Ochsner, 2015). Participants passively viewed or effortfully regulated emotional responses to a series of developmentally-appropriate negative and neutral images (modeled after the International Affective Picture System (IAPS); Lang, Bradley, & Cuthbert, 2008). Each trial began with an instructional cue (2s duration), 'Look' or 'Far'. The former instructs participants to passively view a subsequent stimulus whereas the latter prompts participants to regulate it via a psychological distancing variant of cognitive reappraisal. Following the cue, participants viewed a neutral or aversive stimulus. Aversive stimuli were either paired with 'Look' or 'Far'. Neutral images were always paired with the 'Look' cue; these images are not of interest to the current study and will not be discussed further. Stimuli were displayed for 6-10s (jittered). A 5-point

Likert scale was presented thereafter to collect participants affective responses (4s). Lastly, a jittered ITI concluded each trial (1.5-6.5s). Prior to scanning, participants were trained extensively by experimenters to ensure they understood the task and were comfortable completing it, and afterwards completed a practice run. Stimuli were purchased from a stock photography website (https://shutterstock.com) and were normalized in a sample of youth (120, ages 6-16, 50% female) based on valence, arousal, and dominance. Stimuli and normed ratings are available online (osf.io/43hfq/). The most negatively rated images were selected for use in the current study. Participants completed four runs, each lasting approximately 220s (rungs ranged between 110-115 4D volumes). Each task condition ('Look' – Negative; 'Far' – Negative; 'Look' – Neutral) was presented 20 times over the four runs (5 per run). The task was programmed and presented in E-Prime (Psychology Software Tools, Inc., http://www.pstnet.com). A visual representation of the task is depicted in Figure S1.

fMRI Data Acquisition. If participants were younger than 12 years of age or exhibited any nervousness about scanning, they were taken to a mock MRI scanner in order to familiarize with them with scanning environment and to be trained on how to minimize head motion in the scanner. Additionally, in preparation for the scan, participants were packed into the head coil with an inflated, head-stabilizing pillow to restrict movement.

Images were acquired at the University of Washington's (Seattle) Integrated Brain Imaging Center on a 3T Phillips Achieva scanner. A 32-channel head coil was used in conjunction with parallel image acquisition. A T1-weighted, magnetization-prepared rapid acquisition gradient echo (MPRAGE) image was acquired for registration purposes (TR=2530ms, TE=1640-7040μs, 7º flip angle, 256 mm^2 FOV, 176 slices, 1 mm^3 isotropic voxels). Blood oxygenation level dependent (BOLD) signal during functional runs was recorded using a gradient-echo T2*-weighted echoplanar imaging (EPI) sequence. Thirty-two 3-mm thick slices were acquired parallel to the AC-PC line (TR=2000ms, TE=30ms, 90º flip angle, 256 x 256 FOV, 64 x 64 matrix size).

fMRI Data Pre-Processing. The FMRIB Software Library package (FSL, version 5.0.9; fsl.fmrib.ox.ac.uk) was used for pre-processing and later analysis. Prior to pre-processing, data were visually inspected for artifacts, anatomical abnormalities, and excessive head motion. Pre-processing began by using the brain extraction tool (BET) to remove non-brain tissue from images and estimating the extent of participant head motion by using FSL Motion Outliers to record volumes that exceed a 0.9 mm framewise displacement (FD) threshold (Siegel et al., 2014). Fortunately, head motion was relatively low in this sample. The average number of volumes exceeding our FD threshold per run, per participant was 3.014 (*SD* = 6.09, range = 0 - 30.75). Head motion as not appreciably related with age or emotion regulation ability scores (both *rs* < |.05|). Afterwards, we corrected for head motion by spatially realigning volumes using MCFLIRT and then hi-pass filtered the data (100s cutoff). We pre-whitened the data to correct for temporally autocorrelated residuals.

Network Definition. As previously mentioned, we initially started with four prospective network parcellations to choose from: the Schaefer 7 network parcellation (100 ROIs), the Schaefer 17 network parcellation (400 ROIs), and two NeuroSynth parcellations (Schaefer et al.,

2018). The two NeuroSynth parcellations were based on a synthesis of the literature. The first parcellation (which we call 'NeuroSynth1') was comprised of networks that are theoretically necessary to support cognitive reappraisal: the ventral attention network, the dorsal attention network, the salience network, the default mode network, and the frontoparietal network (Ochsner et al., 2012). The second parcellation (which we call 'NeuroSynth2') was based on brain networks that are theoretically relevant for general affective neurodevelopment: the ventral attention network, the salience network, the default mode network, the frontoparietal network, the somatomotor network, the auditory network, the visual network, and the limbic network (Lee, Miernicki, & Telzer, 2017; Lee & Telzer, 2016). As previously noted, the cognitive reappraisal network (CRN; Buhle et al., 2014) was included in all parcellations (both neurosynth and Schaefer parcellations). For the NeuroSynth parcellations, ROIs of each network were defined in a two-step process. We (i) downloaded meta-analytic maps for each network and then (ii) extracted local peaks from the clusters in each image to serve as our ROIs. The meta-analytic map for the CRN was obtained from a prior meta-analysis (Buhle et al., 2014). For all parcellation, spheres (4 mm radius) were drawn across all peaks in order to extract data necessary to compute connectivity (described in the next section). A 4mm radius was used to minimize overlap between spheres.

Estimating Network Connectivity Metrics

     Network Connectivity via Beta-Series. In order to obtain metrics of network-level activity, we first modeled the task data within-subjects using a Least Squares Single (Mumford et al., 2014) design to yield a beta-series (Rissman et al., 2004). For each run within an individual subject, a fixed-effects GLM was created for each 'Far' (regulation of aversive stimuli) and 'Look-Negative' (passive observation of aversive stimuli) 'target trial'. Thus, a separate GLM was created for each individual target trial. Each GLM consisted of a single-trial regressor which modeled that GLM's target trial, a nuisance regressor modeling all other events in the target trial's condition, regressors for the other two conditions—each modeled as they normal would in a standard univariate analysis—and motion parameter regressors. For example, the first 'Far' trial for a participant would receive its own GLM wherein it was modeled in a single regressor; all other 'Far' trials were placed in a separate nuisance regressor while 'Look-Negative' and 'Look-Neutral' trials modeled in their own respective regressors. Afterwards, a new GLM was created for the *second* 'Far' trial, where it was the 'target trial' and modeled in a single-event regressor and the *first* trial is now modeled in the separate nuisance regressor with other trials of that type. FSL's extended motion parameters estimated from MCFLIRT (rotation+translation, derivatives, and squares of each) and additional volumes that exceeded our FD = 0.9 mm threshold were also added as regressors of non-interest to each first-level LSS GLM.

     Following estimation of the LSS GLMs, we used parameter estimates from each trial-specific GLM to create linear contrast images comparing both trial types of interest ('Far' & 'Look-Negative'), respectively, to baseline. We then extracted estimates of activation for both trial types across a series of ROIs nested within six different networks (see below). This yielded two t x p x n arrays (one for each task condition), where entries to each cell represent the average parameter estimate at the t-th trial for the p-th ROI in the n-th network. For every network, beta-

series amongst all its ROIs were correlated amongst each other (Spearman's Rho), yielding two connectivity matrices (one for each trial type) per network. These matrices describe each network's connectivity for a given task condition (i.e., connectivity matrix), and we refer to them as network connectivity profiles (See Figure 1 for an overview).

Network Differentiation

Using Representational Similarity Analysis to Quantify Network Differentiation. The first feature of network activity we chose to examine was differentiation. In order to estimate network differentiation, we compared profiles of network connectivity when youth were engaging in emotion regulation versus when they were simply viewing emotional stimuli using representational similarity analysis (RSA) on our two task-specific connectivity matrices (Kriegeskorte, 2008; Kriegeskorte et al., 2008). Traditionally, RSA is a form of multi-voxel pattern analysis that relies on the presumption that multi-voxel activation patterns reliably contain information about a specific stimulus (Etzel et al., 2013; Kriegeskorte et al., 2009). Comparison of two stimulus patterns (via correlation) in RSA can reveal the extent to which representations of stimuli encode unique or similar information. However, one is not bound to use RSA solely with multivariate patterns. In fact, RSA can be used to compare any two different representations because information about pattern representations is ultimately abstracted away from the modality in which it was initially acquired and transformed forced into a common space. To this point, prior research has used RSA to compare voxelwise patterns to behavioral patterns (Parkinson et al., 2017), behavioral patterns to other behavioral patterns (Brooks & Freeman, 2018; Stolier et al., 2018), and, most importantly for our purposes, representations of brain connectivity patterns (Lee et al., 2017). We built upon and extended these prior implementations of RSA for our purposes in comparing network connectivity states between task conditions.

For each network and task condition, a r x r connectivity matrix was constructed, where r is the number of ROIs in the network. Representational similarity analyses involved vectorizing the off-diagonal elements of each task condition's various connectivity matrices, Fisher transforming them (for variance stabilization), and then correlating the elements (Spearman). This yielded a scalar value that describes the extent to which a given network's connectivity profile differs between task conditions (see Figure 1). An extreme high value indicates that the pattern of connectivity across the entire network remains consistent during 'Far' (Regulate) and 'Look-Negative' (Emotional Baseline) trials; An extreme low value (anticorrelation) indicates a differentiated connectivity profile between the two task states.

Network Community Structure.

Community Number Differences Between Network States. The second feature of network activity we examined was community structure between task conditions (e.g., when regulating versus when passively viewing emotional stimuli) (Sporns, 2010). That is, we examined the extent to which nodes within a network formed communities, or subclusters comprised of densely interconnected nodes within the network. We were specifically interested in quantifying the number of communities within each network as a function of task condition. To achieve this,

we first used the aforementioned connectivity matrices to create undirected, weighted adjacency matrices. An adjacency matrix summarizes a graph—each row and column represent the ROIs (nodes) that comprise the network, while an entry into the off-diagonal elements of the matrix denotes the presence or absence of a connection (edge) between two nodes. It is customary to set an *edge defining threshold*, a value at which all correlation estimates (e.g., correlation between two ROIs' beta-series) at or falling below the threshold are set to zero, and all estimates above the threshold are set to 1 (unweighted edges) or left at their original value (unweighted edges) (Bolt et al., 2017). In our case, we used weighted edges. We varied the edge defining threshold as part of our specification curve analysis.

We then estimated the number of communities in each task condition, for all networks, for all subjects using the R igraph package's walktrap clustering algorithm (Csardi & Nepusz, 2006; Pons & Latapy, 2005) in R (R Core Team, 2014). The walktrap clustering algorithm (Pons & Latapy, 2005) estimates community structure in a graph using random walks—step-wise paces from one node to another on a graph according to random chance. The algorithm operates on the principle that the probability of taking a walk (i.e., moving over) from one node to another depends on the number of shared connections between nodes. This ultimately means that walks are more likely to become 'trapped' inside a community than a non-community due to the dense interconnections between nodes that define community membership. The number of a given network's communities across both task conditions were extracted from the graphs and differenced ('Far' – 'Look-Negative'), resulting in six metrics of community structure, one per network, per participant. A higher value on this metric indicates a greater number of observed communities within a network during emotion regulation, compared to passively viewing emotional stimuli.

Modularity Differences Between Network States. Our final network metric was the difference in modularity between task conditions for a given network. Building upon the concept of community structure, a graph is said to be highly modular if it exhibits relatively dense interconnections within its communities and relatively sparse connections between them. On the other hand, a graph that is low in modularity has nodes that have relatively equal connections within and between their communities. In other words, modularity can be thought to measure the extent to which "cross-talk" exists between the communities within a network (i.e., a ratio between intra- versus inter-community ties). Modularity metrics across both task states in each network across all participants were obtained using the communities identified in the prior section. The same differencing was taken ('Far' – 'Look-Negative') to obtain a metric of modularity differences between task conditions across each participant's connectome. Notably, because our focus is to understand the degree to which information about development and emotion regulation ability is encoded across a set of brain networks, we opted to focus on *within*-network modularity, rather than *between*-network modularity.

Statistical Analysis

Specification Curve Analysis. Specification curve analysis is typically used as a method to quantify the impact of analytic flexibility on scientific results (Simonsohn et al., 2015). The logic is that there are many different defensible ways to analyze and model a dataset with the

purpose of answering a single research question. One can posit any number of different analytic *specifications* and argue for their validity. Crucially, it is known that results may change dramatically between analytic specifications. A specification curve is intended to capture the degree of variability by aggregating results across the entire set of plausible specifications. In this analysis, one runs a dataset through all possible specifications and plots an output of interest (significance test, parameter estimate, fit statistic, etc.). One can then use this curve in a number of ways (model selection, checking robustness, etc.). Here we chose to run a descriptive specification curve along 4 specification dimensions to determine which analytic approaches were going to lend themselves to building the most predictive models of age and emotion regulation ability (defined as having the lowest Root Mean Square Error, RMSE). The four dimensions were: parcellation (Schaefer 7, Schaefer 17, NeuroSynth1, NeuroSynth2), edge defining threshold ($\rho = 0.4, 0.5, 0.6$), linear de-trending of beta-series via OLS regression (de-trend, do not de-trend) and inclusion of non-brain predictors into the model (including vs excluding predictors). The end goal was to select an optimal modeling strategy among each of the 4 specification dimensions. To do this, we built specification curves for predicting age and emotion regulation ability and selected the 5 best (corresponding to the top ~10%) models. If models conflicted in their specifications, we aggregated over results from models with each specification.

Model Building Using Ridge Regression. We used ridge regression to estimate associations between network activity metrics and facets of development (age, emotion regulation ability). Ridge regression was selected because (i) we had continuous dependent variables, (ii) of its ability to effectively estimate parameters in a model with many predictors, (iii) of its ability to handle highly correlated predictor variables, and (iv) most importantly, its decreased susceptibility to overfitting and therefore better out of sample predictive accuracy (McNeish, 2015; Murphy, 2012). Ridge regression differs from traditional OLS regression in that its loss function contains an extra penalty term (Penalty $b_{Ridge} = \lambda\Sigma b_j^2$). This has the effect of biasing, or shrinking, coefficients towards zero, especially those with inappropriately large magnitudes that drive overfitting. This added bias, introduced via what is known as *l2 regularization*, helps prevent overfitting by decreasing sample to sample variability in regression coefficients. This has the effect of producing more predictive models but also helping make inferences about population parameters due to enhanced certainty in parameter estimates (i.e., narrower sampling distribution of coefficient estimates).

Here we implemented ridge regression using the `glmnet()` function in R. Notably, ridge regression requires a tuning parameter (λ) that controls the degree of regularization. In our case, we used the `cv.glmnet()` function, which finds the optimal λ value via cross-validation. To facilitate statistical inference, we computed 95% bootstrapped confidence intervals around the predicted values for each of the top models from the specification curve results (one set for predicting age, another for predicting emotion regulation ability). We combined results from the top models of our specification curves by averaging over coefficients and 95% CIs.

Taxonomizing and Validating Brain Networks Via Hierarchical Clustering and Bootstrapping. We used agglomerative clustering—implemented with R's `hclust()` function—

to create a working taxonomy of network activity during emotion regulation (Tversky & Hutchinson, 1986). Agglomerative clustering is a 'bottom-up' approach in which the algorithm begins by assigning each item to its own cluster and then iteratively groups items into larger, hierarchical clusters based on similarity (based on distance in Euclidean space). The output of an agglomerative clustering analysis is a dendrogram – the mathematical object that stores information about hierarchical relationships between its elements. This object serves as our working taxonomy of network organization during emotion regulation. We used principle components analysis to combine information from the three metrics of network activity into a single score. Specifically, we extracted principle component scores of the first component to serve as our composite metric of network activity. Finally, to determine whether the stability of the taxonomy (or, if it simply occurred just to chance) we used cophenetic correlation—a method to compare the similarity of dendrograms (Liu, Zhu, Qiu, & Chen, 2012)—in conjunction with bootstrapping. The logic here was that we could resample observations from our data with replacement, build a new dendrogram from the bootstrapped sample, and correlate the bootstrapped dendrogram with our observed estimate. By aggregating cophenetic correlations between the observed and bootstrapped dendrograms, we are effectively building a sampling distribution of the dendrograms and determining how well our observed dendrogram matches those in the sampling distribution. If the distribution of cophenetic correlations contains many values at or below zero, then that would imply our dendrogram comes from a distribution with no stable or consistent structure. This procedure is consistent with related analyses in developmental affective neuroscience (Zhang et al., 2019).

Results

Analytic Overview. In the current study, we scanned 70 children and adolescents with fMRI while they completed four runs of a cognitive reappraisal task that involved alternately regulating (regulate condition) and reacting naturally (emotional baseline condition) to a series of negative and neutral images. Cognitive reappraisal—reframing a stimulus to alter its emotional impact—was our chosen form of emotion regulation study because it is widely-studied, is highly consequential for adjustment (e.g., Troy et al., 2010), its abilities are known to improve with age during childhood and adolescence (Gross, 1998; McRae et al., 2012; Silvers et al., 2012). Before proceeding, we are compelled to note our questions were exploratory in nature. Therefore, our study was guided by open questions about (i) how network properties differentially encode chronological age and emotion regulation ability, and (ii) the taxonomical organization of brain networks involved in the neurodevelopment of emotion regulation rather than formal hypotheses. Accordingly, due to this exploratory nature, we do not correct our analyses for Type I error rate.

In order to predict age and emotion regulation ability from whole-brain network properties, as well as to examine the taxonomical organization of brain networks involved in emotion regulation, it was necessary to first define brain networks. The non-trivial task of parcellating the human brain into functionally meaningful units organized according to some kind of hierarchy is both a new and old problem in neuroscience. While this practice initially began centuries ago, the latest attempts at doing such are no less prone to stirring controversy (Eickhoff, Constable, & Yeo, 2018). Modern parcellation attempts involve aggregating resting state data from thousands of research participants and performing complex clustering analyses

that assign neighboring voxels to individual parcels while delineating boundaries between said parcels, and then assign parcels to whole-brain networks (Schaefer et al., 2018). While this approach has been successful for a number of research purposes (Sripada et al., 2020), it has also been shown that such parcellations are somewhat volatile in the face of task demands (Salehi et al., 2019). To best navigate this conundrum, we conducted a descriptive specification-curve analysis—running and aggregating results of analyses across all plausible analytic specifications—as to determine which parcellation (among 4) simultaneously yielded the most predictive model while also being robust to other analytic choices (Orben & Przybylski, 2019; Simonsohn, Simmons, & Nelson, 2015). The goal of this analysis was to identify the most predictive and robust parcellation scheme among a set of four potential parcellations. Other such analytic choices factored into the specification curve analysis were: (i) whether to detrend each region of interest's (ROI) beta-series (Shumway & Stoffer, 2017), (ii) what edge-defining threshold to use to create adjacency matrices for graph theory analyses, and (iii) whether to include non-brain predictors (e.g., include age when predicting emotion regulation ability and vice versa). More details about this approach is described below in the *specification curve* section of the methods.

The following procedure was conducted to derive the network metrics we would use for both aims articulated above. Following visual and statistical assessment of data quality and preprocessing, we prepared the data for analysis by extracting beta-series from each node (i.e., ROI) across all networks within a parcellation. Various metrics from each network were next calculated with the goal of using said metrics to build classifiers that could predict age and emotion regulation ability Correlations among each networks' beta-series were performed and used, in conjunction with representational similarity analysis (RSA) and graph theory, to yield our metrics of network activity. These metrics included network differentiation between task states (regulate v. emotional baseline; RSA based) and two measures of network community structure: 1) differences in the *number of modules* within a network between task states (regulate v. emotional baseline), and 2) differences in *modularity* between task states (the degree of 'crosstalk' between modules for regulate v. emotional baseline). Notably, the number of modules and modularity are related but distinct—the former tells us how many 'divisions of labor' occur in a network and the latter reflects how rigidly the network appears to follow such divisions. We reiterate here that our inclusion of, and comparison to, a control conditional (a passive emotional baseline) is critical for making inferences about emotion regulation rather than broader emotional processes.

Afterwards we submitted connectivity-based metrics from the top specifications to ridge regression analyses predicting age and emotion regulation ability. As detailed in the methods section, ridge regression is a regression-based method that incorporates small, controlled systematic bias to reduce uncertainty in parameter estimates. We then computed and aggregated the root mean square error (RMSE) as our metric of fit per each specification (taking into account the number of predictors for each model, i.e., waiting by degrees of freedom). We then determined the best parcellation to use based on the specification curve (Schaefer 17), and computed parameter estimates of each network predictor and bootstrapped confidence intervals for statistical inference (across two separate models, one predicting age and another predicting emotion regulation ability). Finally, using the same parcellation, we attempted to establish a working network taxonomy for emotion regulation in youth using our connectivity-based metrics

of brain activity in conjunction with methods typically used in evolutionary biology, namely agglomerative clustering (algorithmic grouping based on shared and unique characteristics) and cophenetic correlation analysis (statistical comparison of clustering solutions). The Methods section presents greater detail on each of these steps. Our results follow.

Behavioral Results. Behavioral results from our emotion regulation task have been reported elsewhere (see Guassi Moreira et al., 2018), but we briefly recapitulate them here for convenience. Participants' average negative affective ratings during the 'Far' and 'Look-Negative' trials were 1.95 (SD = .499) and 2.48 (SD = .562), respectively. This difference was statistically significant ($t(69)$ = -10.37, $p < .001$). Participants' ability to engage in emotion regulation, quantified as the percent reduction in negative affect for the regulate condition relative to the emotional baseline condition, improved with age ($r = 0.344$, $p<.01$).

Using network approaches to parse chronological age and emotion regulation ability

Specification Curve Results.

We subsequently specified four potential network parcellations: Two using the recently published Schaefer atlas (7 networks, 100 ROIs; 17 networks 400 ROIs) and two using a theoretically motivated set of networks whose ROIs were defined via NeuroSynth (Schaefer et al., 2018). The two network sets from the Schaefer atlas were derived from resting state scans, whereas the network sets from NeuroSynth were defined by meta-analyzing coordinates from both resting state and task-based fMRI studies. We also included a 'cognitive reappraisal network' (commonly observed during cognitive reappraisal) to all parcellations defined via meta-analysis (CRN; Buhle et al., 2014) because it is the most relevant network to the emotion regulation task incorporated in the present study. This is important because recent work in both adult and developmental samples suggests that task-specific information—and even activity from task-specific networks—is privileged over domain-general information and thus might be highly informative when differentiating between chronological age and task ability (Brody et al., 2019; Greene, Gao, Scheinost, & Constable, 2018).

Results for specification curves for both emotion regulation ability and age are visualized in Figure 2 and the top five models (~top 10%) for each are summarized in Table 1. While both sets of curves contained highly predictive models, we refrain from over-interpreting RMSE here because (i) RMSE was used as the selection criteria to determine our analytic specifications and (ii) our stated goal was to dissociate relationships between brain network activity and developmental constructs of regulation and ability and general maturation (i.e., age). Nevertheless, we do note for exploratory purposes that both the top emotion regulation ability models were all under 1, whereas the top models for age ranged just below and just above 1. Descriptively this suggests that, in our sample, brain connectivity during emotion regulation is a better predictor of regulatory ability than age. Importantly, both curves converged on the same parcellation, the 17 network Schaefer parcellation. Therefore, all subsequent predictive models of age and emotion regulation ability were implemented using the 17 network Schaefer parcellation, which is comprised of 7 supranetworks (control network, default mode network, salience-ventral attention network, dorsal attention, visual network, limbic network, somatomotor network) and their subnetworks, in addition to our inclusion of the CRN. By

contrast, specification curves did not reveal a clear consensus on the best edge defining threshold, detrending strategy, or strategy for controlling for non-brain variables. To address this issue, we ran our analyses for age and ability five times–one for each of the top (respective) specifications—and averaged the coefficients across the set of analyses. Results for each dependent variable (age, emotion regulation ability) are reported below. Bootstrapped 95% confidence intervals (5,000 iterations) were computed in order to perform inference over the parameter estimates.

Age. None of the network connectivity metrics successfully predicted age above chance. Emotion regulation ability was a significant predictor of age ($b = 0.0279$, 95% bootstrapped CI = [0.0479, 0.3523]), even after controlling for network connectivity metrics. No network connectivity metrics emerged as significant predictors when removing emotion regulation ability from the model.

Emotion Regulation Ability. When predicting emotion regulation ability from our network connectivity metrics, we found that network differentiation in Control Network (subnetwork C) and network modularity in the Default Mode Network (subnetwork B) were each predictive of emotion regulation ability. Differentiation in the control network was related with emotion regulation ability ($b = -0.0640$, 95% bootstrapped CI = [-0.3136, -0.0197]), such that greater control network differentiation was related to better emotion regulation ability (note that greater differentiation is indexed by relatively lower (more anticorrelated) network connectivity profiles, explaining while the sign here is negative). Modularity in the DMN was directly related with emotion regulation ability ($b = 0.0652$, 95% bootstrapped CI = [0.0240, 0.2784]), such that greater modularity during reappraisal relative to an emotional baseline was associated with better emotion regulation ability. These effects were significant regardless whether age was included as a covariate in the model. Age was also a significant predictor of emotion regulation ability in this model ($b = 0.0817$, 95% bootstrapped CI = [0.0593, 0.3772]), even after accounting for network activity metrics. A follow-up analysis indicated that age did not interact with either control network differentiation or default mode network modularity to predict emotion regulation ability. Last, in order to determine the variance in ability that control network differentiation, default mode modularity, and age each uniquely predicted, we calculated squared semi-partial correlations. Control network differentiation uniquely accounted for 6.866% of the variance in emotion regulation ability, default mode network modularity uniquely accounted for 10.298% of such variance, and age uniquely accounted for 7.295% of such variance. Combined, the two network connectivity metrics uniquely accounted for 16.590% of the variance in emotion regulation ability scores (relative to age).

Creating and Testing Taxonomies of Emotion Regulation Neurodevelopment

The final aim was to provide a neurodevelopmental taxonomy of brain networks implicated in emotion regulation. The section below describes the creation and validation of three working network taxonomies.

Taxonomy Building and Validation. We used principal components analysis (PCA) and agglomerative clustering to build our network taxonomy (Figure 3). Just as with the prior section

on model building, we used the Schaefer 17 network parcellation. As described in the methods, PCA was performed to reduce the data from three metrics into a single composite, and then a clustering algorithm was used to group networks based on shared network properties. A glance at the dendrogram reveals several interesting features, beginning with the fact that four supraordinate clusters of networks emerged (indicated by color in Figure 3). The largest cluster included the CRN and several other networks implicated in perceptual processes (e.g., dorsal attention network) and higher-order cognition (control network). The limbic network appears to be taxonomically similar to the default mode network, and higher order cognitive networks tasked with cognitive and control and attention were taxonomically related to somatomotor networks.

To test whether the organization of the dendrogram was due to chance, we bootstrapped our sample, re-created the dendrogram at each iteration, and correlated the bootstrapped dendrogram at a given iteration with our observed dendrograms (The cor_cophenetic()function of R's 'dendextend' package was used to compute cophenetic correlations, a metric that compares similarity of two dendrograms). The ensuing distribution of correlations (c(observed, a bootstrapping iteration) constituted our empirical sampling distribution. If the 95% confidence intervals were both greater than 0, we rejected the notion that the dendrogram's organization capitalized purely on noise. We also considered the proportion of bootstrapped cophenetic correlations that were equal to or below the null value of zero to account for skewed bootstrapped distributions. More detail about our inferential procedures and logic is given in the methods. Results from bootstrapping suggest that the structure of our dendrogram was unlikely to have occurred by chance ($c = 0.2671$, 95% bootstrapped CI [0.0111,0.5864], bootstrapped cs $\leq$ 0: 2.19%).

Discussion

We set out to interrogate the neurodevelopment of emotion regulation using network neuroscience. By doing so, our goal was to parse the brain networks that contribute to correlated but distinct developmental processes (chronological age and acquisition of emotion regulation abilities), as well as establish a working neurodevelopmental taxonomy of brain network activity during emotion regulation. This work builds upon past network neuroscience studies of development to better understand task-specific aspects of maturation (Rudolph et al., 2017). We found that emotion regulation ability showed significant associations with whole-brain network metrics in the control network and default mode network. Because greater network differentiation and modularity in these networks was associated with better regulation ability, it suggests these networks develop specialized features that lend themselves to better implementing emotion regulation. Importantly, these combined network properties explained a substantial amount of the variance in emotion regulation ability over and above age. Taxonomical analyses revealed that that the CRN was similar to networks implicated in higher-order cognitive processes as well as perception, while limbic and default mode networks evinced similar taxonomic organization to one another. These findings have implications for our basic understanding of the neurodevelopment of emotion regulation, and how brain networks may broadly change and interact across cognitive and emotional development.

Brain Networks Encode Regulatory Ability, but not Age. In the present study, we were able to predict emotion regulation ability from network activity metrics, even after accounting for age. That our models were unable to predict age—despite its strong association with emotion regulation ability—speaks to their specificity. These results suggest that task-based network activity encodes information about skill development associated with a given task rather than domain general features of development (e.g., age). We found that greater differentiation in the control network, and greater modularity in the default mode network, were each associated with better emotion regulation ability. Prior work has found that increasing structural modularity of the control network is associated with increased executive control skills across age (Baum et al., 2017), suggesting that the anatomical substrates of the network become primed to support functional specialization. Our study builds upon this result by showing that individuals with highly distinct connectivity states during regulation relative to a passive emotional baseline are better able to regulate their emotions, even while controlling for age. Critically, we are able to make this inference because reappraisal is a cognitively demanding emotion regulation strategy that mechanistically relies on a suite of executive functions (Etkin, Büchel, & Gross, 2015; Ochsner, Silvers, & Buhle, 2012). We found a conceptually similar result with the default mode network, as individuals whose default mode networks exhibited greater modularity during regulation compared to an emotional baseline were also more likely to possess better emotion regulation skills. This work suggests that divisions of labor within the default mode network become more pronounced with skill acquisition, consistent with recent findings in another cognitive domain showing that increasing subdivision of the default mode network facilitated better mentalizing abilities on a theory of mind task (Richardson et al., 2018). Because reappraisal requires self-monitoring and introspection (Ochsner et al., 2012), this increasing modularity in default mode network could reflect increasingly mature mentalizing processes.

Prior work has largely been able to predict age from resting state data yet has been less successful at doing so from task data (Dosenbach et al., 2010), consistent with our null age results. This prior finding and our present results point to the possibility that resting brain activity reflects domain-general information about development, whereas task-related activity better reflects domain-specific skills (Greene et al., 2018; Kaufmann et al., 2017). This speaks to the nature of neurodevelopment, suggesting that maturation is driven by changes in network states, unique and discrete configurations of activity and connectivity that a single network could adopt at any time (Cornblath et al., 2018). This fact implies network states may emerge depending on a network's position along a limited number of unobserved axes (Shine, et al., 2019ab). Neurodevelopment may be subsequently conceptualized as a network's trajectory over time in this space, with movement along certain dimensions more consequential for general maturation, and movement along others more impactful for shaping task-specific development. If true, it would appear brain networks in our data were only configuring themselves in states that are sensitive to regulation ability but not age. Our data preclude us from verifying this possibility, but our results, in conjunctions with other, imply what may be a fruitful line of research.

Our results also stand to inform models of the neural bases of emotion regulation. Current accounts of emotion regulation posit that individual differences in emotion regulation ability are driven by activity among a set of brain regions that comprise a dedicated network for a given

regulatory strategy (Ochsner et al., 2012). Such research has only recently begun to consider whole-brain networks as meaningful biological units (Zhang et al., 2019), yet our results highlight the importance of studying whole-brain networks in the context of emotion regulation. Although the notion of dedicated networks for emotion regulation strategy is enticing, our research fails to find support that a cognitive reappraisal network is associated with individual differences in emotion regulation ability. This may be because the network itself does not constitute a meaningful biological entity despite prior evidence indicating otherwise or because variability in both network activity and regulation ability were truncated from an otherwise full range in our sample (i.e., we only sampled a limited range of the population values of network activity metrics). Yet another possibility is that CRN activity during reappraisal may actually emerge as an extension or special form of another network whose core activity might be more tightly linked with regulatory ability.

A (Working) Neurodevelopmental Taxonomy of Brain Network States During Emotion Regulation. The results from our taxonomy building analyses suggest that brain network activity during active regulation in youth evinces a coherent structure. That the limbic network and default mode network were consistently grouped together implies shared activity profiles between the networks, which in turn signify the possibility that self-referential processes during emotion regulation are consequential for shaping bottom-up affective responses and conforms to long-held theories about emotion regulation (Ochsner et al., 2012). Likewise, that the CRN and a suite of networks engaged in higher-order cognitive and perceptual processes were repeatedly grouped together may suggest that the CRN is actually a specific and emergent state among these networks. This fits with popular models of the cognitive control of emotion, which suggest that executive processes may be similarly deployed in affective and non-affective contexts (Ochsner & Gross, 2005). Alternately, it is also possible the CRN acts as a manager or integrator of multiple cognitive networks that must implement the cognitive mechanics of reappraisal. Such a notion has been posited with respect to the functioning of individual ROIs (Ligeza et al., 2016; Morawetz et al., 2016), but never in the context of whole-brain networks. Last, subdivisions of the control network were variable in their taxonomic relationship with other networks, likely showcasing this network's flexibility in shaping brain activity to accommodate task demands (Marek & Dosenbach, 2018). These results from our working taxonomy have several potential applications for future research. Notably, understanding how networks are typically organized across childhood and adolescence may help to establish standards of healthy brain development by providing an organizational standard against which individuals can be compared. Such efforts could ultimately facilitate deep phenotyping and precision medicine applications in pediatric neuroscience (Robinson, 2012).

Limitations and Future Directions. This study is accompanied by several limitations. First, although our sample size is rather large for pediatric neuroimaging, and is twice the median cell size of human fMRI studies (Poldrack et al., 2017), it is nevertheless somewhat small in comparison to other network neuroscience studies (e.g., Baum et al., 2017; DuPre & Spreng, 2017). Further, although cross-sectional research is an adequate starting point for characterizing the neurodevelopment of emotion regulation, future research would benefit in having longitudinal data to examine how the trajectory of each network's activity change across time,

helping us better clarify the development of emotion regulation. Additionally, amid growing concerns about reliability (Elliott et al., 2019; Enkavi et al., 2019), we feel compelled to note that the behavioral and neural test-retest reliability of the emotion regulation task used here has not yet been established. A final future direction involves testing whether the predictive properties of the network studied here hold in other emotional states, or even across different types of emotion regulation (e.g., suppression). Addressing these concerns under further empirical scrutiny will help lend further confidence in the take-away conclusions raised by our results.

In sum, this study is the first to our knowledge that examines how the developing brain supports the acquisition of emotion regulation—a cornerstone of mental health—independent of age. We show that age and emotion regulation abilities are separately encoded in patterns of network activity, while also highlighting that these networks fit into a broader taxonomic structure.

References

Baum, G. L., Ciric, R., Roalf, D. R., Betzel, R. F., Moore, T. M., Shinohara, R. T., … Satterthwaite, T. D. (2017). Modular segregation of structural brain networks supports the development of executive function. *Current Biology*, *27*(11), 1561-1572.e8. https://doi.org/10.1016/j.cub.2017.04.051

Bolt, T., Nomi, J. S., Yeo, B. T. T., & Uddin, L. Q. (2017). Data-driven extraction of a nested model of human brain function. *Journal of Neuroscience*, *0323–17*. https://doi.org/10.1523/JNEUROSCI.0323-17.2017

Braunstein, L. M., Gross, J. J., & Ochsner, K. N. (2017). Explicit and implicit emotion regulation: A multi-level framework. *Social Cognitive and Affective Neuroscience*, *ePub ahead*, 1–13. https://doi.org/10.1080/02699931.2010.544160

Brody, G. H., Yu, T., Nusslock, R., Barton, A. W., Miller, G. E., Chen, E., … Sweet, L. H. (2019). The protective effects of supportive parenting on the relationship between adolescent poverty and resting-state functional brain connectivity during adulthood. *Psychological Science*, 1–9. https://doi.org/10.1177/0956797619847989

Brooks, J. A., & Freeman, J. B. (2018). Conceptual knowledge predicts the representational structure of facial emotion perception. *Nature Human Behaviour*. https://doi.org/10.1038/s41562-018-0376-6

Buhle, J. T., Silvers, J. A., Wage, T. D., Lopez, R., Onyemekwu, C., Kober, H., … Ochsner, K. N. (2014). Cognitive reappraisal of emotion: A meta-analysis of human neuroimaging studies. *Cerebral Cortex*, *24*(11), 2981–2990. https://doi.org/10.1093/cercor/bht154

Calkins, S. (1994). Origins and outcomes of individual differences in emotion regulation. *Monographs of the Society for Research in Child Development*, *59*(2–3), 53–72.

Casey, B. J., Jones, R. M., & Hare, T. A. (2008). The adolescent brain. *Annals of the New York Academy of Sciences*, *1124*, 111–126. https://doi.org/10.1196/annals.1440.010

Cornblath, E. J., Ashourvan, A., Kim, J. Z., Betzel, R. F., Ciric, R., Baum, L., … Bassett, D. S. (2018). Context-dependent architecture of brain state dynamics is explained by white matter connectivity and theories of network control. *BioRxiv*, 1–33.

Csardi, G., & Nepusz, T. (2006). The igraph software package for complex network research. *InterJournal*, *Complex Sy*, 1695.

Dosenbach, N. U. F., Nardos, B., Cohen, A. L., Fair, D. a, Power, J. D., Church, J. a, … Schlaggar, B. L. (2010). Prediction of individual brain maturity using fMRI. *Science*, *329*(5997), 1358–1361. https://doi.org/10.1126/science.1194144

DuPre, E., & Spreng, R. N. (2017). Structural covariance networks across the lifespan, from 6-94 years of age. *Network Neuroscience*, 1–38. https://doi.org/10.1162/NETN_a_00016

Eickhoff, S. B., Constable, R. T., & Yeo, B. T. T. (2018). Topographic organization of the cerebral cortex and brain cartography. *NeuroImage*, *170*, 332–347. https://doi.org/10.1016/j.neuroimage.2017.02.018

Eisenberg, I., Bissett, P., Enkavi, A. Z., Li, J., MacKinnon, D., Marsch, L., & Poldrack, R. (2019). Uncovering mental structure through data-driven ontology discovery. *Nature Communications*, *10*(1), 2319–2332. https://doi.org/10.31234/OSF.IO/FVQEJ

Elliott, M. L., Knodt, A. R., Ireland, D., Morris, M. L., Ramrakha, S., Sison, M. L., … Hariri, A. R. (2019). Poor test-retest reliability of task-fMRI: New empirical evidence and a meta-analysis. *BioRxiv*, (919), 1–31.

Enkavi, A. Z., Eisenberg, I. W., Bissett, P. G., Mazza, G. L., Mackinnon, D. P., Marsch, L. A., & Poldrack, R. A. (2019). Large-scale analysis of test – retest reliabilities of self-regulation measures. *Proceedings of the National Academy of Sciences*, *116*(12), 5472–5477. https://doi.org/10.1073/pnas.1818430116

Etkin, A., Büchel, C., & Gross, J. J. (2015). The neural bases of emotion regulation. *Nature Reviews Neuroscience*, *16*(11), 693–700. https://doi.org/10.1038/nrn4044

Etzel, J. A., Zacks, J. M., & Braver, T. S. (2013). Searchlight analysis: Promise, pitfalls, and potential. *NeuroImage*, *78*, 261–269. https://doi.org/10.1016/j.actbio.2009.04.013.Role

Fair, D. A., Cohen, A. L., Dosenbach, N. U. F., Church, J. A., Miezin, F. M., Barch, D. M., … Schlaggar, B. L. (2008). The maturing architecture of the brain's default network. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(10), 4028–4032. https://doi.org/10.1073/pnas.0800376105

Gratton, C., Laumann, T. O., Nielsen, A. N., Greene, D. J., Gordon, E. M., Gilmore, A. W., … Petersen, S. E. (2018). Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. *Neuron*, *98*(2), 439-452.e5. https://doi.org/10.1016/j.neuron.2018.03.035

Greene, A. S., Gao, S., Scheinost, D., & Constable, R. T. (2018). Task-induced brain state manipulation improves prediction of individual traits. *Nature Communications*, *9*(1), 1–13. https://doi.org/10.1038/s41467-018-04920-3

Gross, J. J. (1998). The emerging field of emotion regulation: an integrative review. *Review of General Psychology*, *2*(5), 271–299. https://doi.org/10.1017.S0048577201393198

Gross, J. J. (2015). Emotion regulation: Current status and future prospects. *Psychological Inquiry*, *26*(1), 1–26. https://doi.org/10.1080/1047840X.2014.940781

Guassi Moreira, J. F., McLaughlin, K. A., & Silvers, J. A. (2019). Spatial and temporal cortical variability track with age and affective experience during emotion regulation in youth. *Developmental Psychology*, *55*(9), 1921–1937.

Guest, O., & Love, B. C. (2017). What the success of brain imaging implies about the neural code. *ELife*, *6*, 1–16. https://doi.org/10.7554/eLife.21397

Kaufmann, T., Alnæs, D., Doan, N. T., Brandt, C. L., Andreassen, O. A., & Westlye, L. T. (2017). Delayed stabilization and individualization in connectome development are related to psychiatric disorders. *Nature Neuroscience*, *20*(4), 513–515. https://doi.org/10.1038/nn.4511

Kleckner, I. R., Zhang, J., Touroutoglou, A., Chanes, L., Xia, C., Simmons, W. K., … Feldman

Barrett, L. (2017). Evidence for a large-scale brain system supporting allostasis and interoception in humans. *Nature Human Behaviour*, *1*(5). https://doi.org/10.1038/s41562-017-0069

Kriegeskorte, N. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *2*, 1–28. https://doi.org/10.3389/neuro.06.004.2008

Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., … Bandettini, P. A. (2008). Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, *60*(6), 1126–1141. https://doi.org/10.1016/j.neuron.2008.10.043

Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S., & Baker, C. I. (2009). Circular analysis in systems neuroscience: The dangers of double dipping. *Nature Neuroscience*, *12*(5), 535–540. https://doi.org/10.1038/nn.2303

Lang, P. J., Bradley, M. M., & Cuthbert, B. N. (2008). *International affective picture system (IAPS): Affective ratings of pictures and instruction manual. Technical Report A-8.* https://doi.org/10.1016/j.epsr.2006.03.016

Lee, T. H., Miernicki, M. E., & Telzer, E. H. (2017). Families that fire together smile together: Resting state connectome similarity and daily emotional synchrony in parent-child dyads. *NeuroImage*, *152*, 31–37. https://doi.org/10.1016/j.neuroimage.2017.02.078

Lee, T. H., & Telzer, E. H. (2016). Negative functional coupling between the right fronto-parietal and limbic resting state networks predicts increased self-control and later substance use onset in adolescence. *Developmental Cognitive Neuroscience*, *20*, 35–42. https://doi.org/10.1016/j.dcn.2016.06.002

Ligeza, T. S., Wyczesany, M., Tymorek, A. D., & Kamiński, M. (2016). Interactions between the prefrontal cortex and attentional systems during volitional affective regulation: An effective connectivity reappraisal study. *Brain Topography*, *29*(2), 253–261. https://doi.org/10.1007/s10548-015-0454-2

Linnæus, C. (1758). *Systema Naturae*. Gmelin.

Liu, X., Zhu, X.-H., Qiu, P., & Chen, W. (2012). A correlation-matrix-based hierarchical clustering method for functional connectivity analysis. *Journal of Neuroscience Methods*, *211*(1), 94–102. https://doi.org/10.1016/j.jneumeth.2012.08.016.A

Marek, S., & Dosenbach, N. U. F. (2018). The frontoparietal network: function, electrophysiology, and importance of individual precision mapping. *Dialogues in Clinical Neuroscience*, *20*, 133–140. Retrieved from www.dialogues-cns.org

McNeish, D. M. (2015). Using lasso for predictor selection and to assuage overfitting: A method long overlooked in behavioral sciences. *Multivariate Behavioral Research*, *50*, 471–483. https://doi.org/10.1080/00273171.2015.1036965

McRae, K., Gross, J. J., Weber, J., Robertson, E. R., Sokol-Hessner, P., Ray, R. D., … Ochsner, K. N. (2012). The development of emotion regulation: An fMRI study of cognitive reappraisal in children, adolescents and young adults. *Social Cognitive and Affective Neuroscience*, *7*(1), 11–22. https://doi.org/10.1093/scan/nsr093

Morawetz, C., Bode, S., Baudewig, J., Kirilina, E., & Heekeren, H. R. (2016). Changes in Effective Connectivity Between Dorsal and Ventral Prefrontal Regions Moderate Emotion Regulation. *Cerebral Cortex*, *26*(5), 1923–1937. https://doi.org/10.1093/cercor/bhv005

Mumford, J. A., Davis, T., & Poldrack, R. A. (2014). The impact of study design on pattern estimation for single-trial multivariate pattern analysis. *NeuroImage*, *103*, 130–138. https://doi.org/10.1016/j.neuroimage.2014.09.026

Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. Cambridge (Ma): The MIT Press.

Ochsner, K. N., & Gross, J. J. (2005). The cognitive control of emotion. *Trends in Cognitive Sciences*.

Ochsner, K. N., Silvers, J. A., & Buhle, J. T. (2012). Review and evolving model of the cognitive control of emotion. *Annals of the New York Academy of Sciences*, (1251), E1–E24. https://doi.org/10.1111/j.1749-6632.2012.06751.x.Functional

Orben, A., & Przybylski, A. K. (2019). Screens, teens, and psychological well-being: Evidence from three time-use-diary studies. *Psychological Science*, *30*(5), 682–696. https://doi.org/10.1177/0956797619830329

Parkinson, C., Kleinbaum, A. M., & Wheatley, T. (2017). Spontaneous neural encoding of social network position. *Nature Human Behaviour*, *1*(5), 1–7. https://doi.org/10.1038/s41562-017-0072

Poldrack, R. A., Baker, C. I., Durnez, J., Gorgolewski, K. J., Matthews, P. M., Munafò, M., … Yarkoni, T. (2017). Scanning the horizon: Towards transparent and reproducible neuroimaging research. *Nature Reviews Neuroscience*, *18*, 115–126. https://doi.org/http://dx.doi.org/10.1101/059188

Pons, P., & Latapy, M. (2005). Computing communities in large networks using random walks. In *International Symposium on Computer and Information Sciences* (pp. 284–293). Berlin, Heidelberg: Springer.

Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., … Petersen, S. E. (2011). Functional network rrganization of the human brain. *Neuron*, *72*(4), 665–678. https://doi.org/10.1016/j.neuron.2011.09.006

Richardson, H., Lisandrelli, G., Riobueno-Naylor, A., & Saxe, R. (2018). Development of the social brain from age three to twelve years. *Nature Communications*, *9*(1), 1027. https://doi.org/10.1038/s41467-018-03399-2

Rissman, J., Gazzaley, A., & D'Esposito, M. (2004). Measuring functional connectivity during distinct stages of a cognitive task. *NeuroImage*, *23*(2), 752–763. https://doi.org/10.1016/j.neuroimage.2004.06.035

Robinson, P. N. (2012). Deep phenotyping for precision medicine. *Human Mutation*, *33*(5), 777–780. https://doi.org/10.1002/humu.22080

Rudolph, M. D., Miranda-Domínguez, O., Cohen, A. O., Breiner, K., Steinberg, L., Bonnie, R. J., … Fair, D. A. (2017). At risk of being risky: The relationship between "brain age" under

emotional states and risk preference. *Developmental Cognitive Neuroscience*, *24*, 93–106. https://doi.org/10.1016/j.dcn.2017.01.010

Salehi, M., Greene, A. S., Karbasi, A., Shen, X., & Scheinost, D. (2019). There is no single functional atlas even for a single individual: Parcellation of the human brain is state dependent. *BioRxiv*, 1–34.

Sarkheil, P., Zilverstand, A., Kilian-hütten, N., Schneider, F., Goebel, R., & Mathiak, K. (2015). fMRI feedback enhances emotion regulation as evidenced by a reduced amygdala response. *Behavioural Brain Research*, *281*, 326–332. https://doi.org/10.1016/j.bbr.2014.11.027

Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X., Holmes, A. J., … Yeo, B. T. T. (2018). Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral Cortex*, *25*, 3095–3114. https://doi.org/10.1093/cercor/bhx179

Shine, J. M., Breakspear, M., Bell, P. T., Martens, K. A. E., Shine, R., Koyejo, O., … Poldrack, R. A. (2019). Human cognition involves dynamic integration of neural activity and neuromodulatory systems. *Nature Neuroscience*, *22*, 289–299. https://doi.org/10.1038/s41593-018-0312-0

Shine, J. M., Hearne, L. J., Breakspear, M., Poldrack, R. A., Mattingley, J. B., & Cocchi, L. (2019). The low-dimensional neural architecture of cognitive complexity is related to activity in medial thalamic nuclei. *Neuron*, *104*, 1–7. https://doi.org/10.1016/j.neuron.2019.09.002

Shumway, R. H., & Stoffer, D. S. (2017). *Time Series Analysis and its Applications* (4th ed.). Gewerbestrasse, Switzerland: Springer.

Silvers, J. A., McRae, K., Gabrieli, J. D. E., Gross, J. J., Remy, K. A., & Ochsner, K. N. (2012). Age-related differences in emotional reactivity, regulation, and rejection sensitivity in adolescence. *Emotion*, *12*(6), 1235–1247. https://doi.org/10.1037/a0028297.Age-Related

Silvers, J. A., Shu, J., Hubbard, A. D., Weber, J., & Ochsner, K. N. (2015). Concurrent and lasting effects of emotion regulation on amygdala response in adolescence and young adulthood. *Developmental Science*, *18*(5), 771–784. https://doi.org/10.1111/desc.12260

Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2015). Specification curve: Descriptive and inferential statistics on all reasonable specifications, 1–19.

Sizemore, A. E., Giusti, C., Kahn, A., Vettel, J. M., Betzel, R. F., & Bassett, D. S. (2018). Cliques and cavities in the human connectome. *Journal of Computational Neuroscience*, *44*, 115–145.

Sporns, O. (2010). *Networks of the Brain*. MIT Press.

Sripada, C., Taxali, A., Angstadt, M., & Rutherford, S. (2020). Boost in Test-Retest Reliability in Resting State fMRI with Predictive Modeling. *BioRxiv*, 1–25.

Stolier, R. M., Hehman, E., Keller, M. D., Walker, M., & Freeman, J. B. (2018). The conceptual structure of face impressions. *Proceedings of the National Academy of Sciences*. https://doi.org/10.31234/osf.io/d9ha7

Thompson, R., Lewis, M. D., & Calkins, S. (2008). Reassessing Emotion Regulation. *Child Development Perspectives*, *2*(3), 124–131.

Troy, A. S., Wilhelm, F. H., Shallcross, A. J., & Mauss, I. B. (2010). Seeing the silver lining: Cognitive reappraisal ability moderates the relationship between stress and depressive symptoms. *Emotion*, *10*(6), 783–795. https://doi.org/10.1037/a0020262.Seeing

Tversky, A., & Hutchinson, J. W. (1986). Nearest neighbor analysis of psychological spaces. *Psychological Review*, *93*(1), 3–22.

Wagner, W. H. (1950). *The Fern Genus Diellia: Structure, Affinities, and Taxonomy*. University of California, Berkeley.

Yeo, B. T. T., Krienen, F. M., Eickhoff, S. B., Yaakub, S. N., Fox, P. T., & Buckner, R. L. (2015). Functional Specialization and Flexibility in Human Association Cortex. *Cerebral Cortex*, *25*(10), 3654–3672. https://doi.org/10.1093/cercor/bhu217

Yeo, B. T. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., … Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, *106*, 1125–1165. https://doi.org/10.1152/jn.00338.2011.

Zhang, Y., Padmanabhan, A., Gross, J. J., & Menon, V. (2019). Development of human emotion circuits investigated using a big-data analytic approach: Stability, reliability, and robustness. *Journal of Neuroscience*, *39*(36), 7155–7172.

Figure 1. Overview of network activity metrics estimation.

Note. Top panel (Network Connectivity via Beta-Series) depicts the beta-series estimated for each task condition (regulate ('Far'), emotional baseline ('Look-Negative')), for the n-th network. Each column corresponds to the p-th ROI, each row to the t-th trial, and each entry is the average beta-value for a given ROI at a given trial. Beta-series for each subject served as the substrates for estimating our network activity metrics of interest, depicted in the bottom row of panels. The middle panel (Network Differentiation via RSA) depicts process of computing network differentiation. The two boxes represent connectivity profiles (matrices) for each condition for the nth network. Each entry to the matrix is the connectivity estimate between a given pair of ROIs. Unique off-diagonal elements were vectorized and then correlated (Spearman's Rho) to yield a single, scalar value of network connectivity similarity. The bottom panel (Network Community Structure via Graph Theory) depicts graphs for each task condition for the n-th network; each node represents an ROI within a given network, whereas each edge is the connectivity between a pair of ROIs. This step allowed us to compute the number of network communities (i.e., number of modules/communities in each network) and modularity. Network connectivity matrices were thresholded to create adjacency matrices and subsequently graphs. Estimates of community structure and modularity were compared across both task conditions for each network.

Figure 2. Specification curves for predicting age (left) and emotion regulation ability (right) from network activity metrics.

Note. Specification curves for each dependent variable (age, emotion regulation ability) reflect the results of 48 models. Each dot reflects the RMSE (our metric of model fit) for a given model (lower RMSE, better model fit). Each dot corresponds to a unique specification involving brain parcellation (two from NeuroSynth, two from the recently published Schaefer atlas), edge defining threshold for graph theory analyses (.4, .5, .6), whether to linearly detrend beta-series (yes, no), and whether or not to include non-brain predictors (e.g., include ability when predicting age, include age when predicting ability). Each panel is ordered based on RMSE, therefore analysis numbers in both plots do not necessarily correspond to the same specification. Specifications to the five best fitting models for age and emotion regulation ability are detailed in Table 1.

Figure 3. Network taxonomies defined by differentiation, community number and modularity.

Note. This taxonomy was achieved Branches indicate splits in the taxonomy's nested hierarchy. Colors correspond to one of four post-hoc clusters that were identified in the dendrogram. CRN = Cognitive Reappraisal Network; DMN = Default Mode Network (contains A/B/C subdivisions); SaVaN = Salience/Ventral Attention Network (contains A/B subdivisions); CoN = Control Network (contains A/B/C subdivisions); LiN = Limbic Network (contains A/B subdivisions); DAN = Dorsal Attention Network (contains A/B subdivisions); VCN = Visual Center Network; VPN = Visual Peripheral Network; SMN = Somatomotor Network (contains A/B subdivisions); TPN = Temporal Parietal Network. Stability of the taxonomy was inferred using bootstrapping techniques (see methods).

Table 1. Top Specification Curve Results

Rank	Age Specification	Ability Specification
1	Schaefer 17 – 0.4 Thresh – Detrend – Ability	Schaefer 17 – 0.4 Thresh – Detrend – Age
2	Schaefer 17 – 0.4 Thresh – Detrend – Ability	Schaefer 17 – 0.4 Thresh – No-Detrend – Age
3	Schaefer 17 – 0.4 Thresh – No-Detrend – Ability	Schaefer 17 – 0.4 Thresh –Detrend – Age
4	Schaefer 17 – 0.6 Thresh – No-Detrend – Ability	Schaefer 17 – 0.5 Thresh – No-Detrend – Age
5	Schaefer 17 – 0.5 Thresh – Detrend – Ability	Schaefer 17 – 0.6 Thresh – No-Detrend – Age

Note. Ability refers to emotion regulation ability (operationalized as the percent change between mean affective ratings during reappraisal and emotional baseline conditions). Rank reflects better RMSE in descending order (lower numbers reflect lower, therefore better, RMSE values); Schafer 17 refers to the 17 network Schaefer parcellation; 'Thresh' refers to the value of the edge-defining threshold for graph theory analyses; 'Detrend' refers to linear trending of beta-series prior to analysis (via OLS regression), 'No-Detrend' indicates no such preprocessing; the top models for ability included age as a predictor while the top models for age included ability as a predictor.