

Running head: NETWORK RECONFIGURATION DURING LEARNING

Functional brain network reconfiguration during learning in a dynamic environment

Chang-Hao Kao¹, Ankit N. Khambhati², Danielle S. Bassett^{3,4,5,6,7}, Matthew R. Nassar^{8,9}, Joseph T. McGuire¹⁰, Joshua I. Gold¹¹, Joseph W. Kable^{1,*}

¹Department of Psychology, University of Pennsylvania, Philadelphia, PA 19104, USA

²Department of Neurological Surgery, University of California, San Francisco, CA 94122, USA

³Department of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104, USA

⁴Department of Neurology, University of Pennsylvania, Philadelphia, PA 19104, USA

⁵Department of Electrical & Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA

⁶Department of Physics & Astronomy, University of Pennsylvania, Philadelphia, PA 19104 USA

⁷Department of Psychiatry, University of Pennsylvania, Philadelphia, PA 19104, USA

⁸Department of Neuroscience, Brown University, Providence, RI 02912, USA

⁹Robert J. and Nancy D. Carney Institute for Brain Science, Brown University, Providence, RI 02912, USA

¹⁰Department of Psychological and Brain Sciences, Boston University, Boston, MA 02215, USA

¹¹Department of Neuroscience, University of Pennsylvania, Philadelphia, PA 19104, USA

* Correspondence: kable@psych.upenn.edu

Abstract

In dynamic environments, people should learn faster, adjusting their beliefs more in response to new evidence, either when that new evidence is more surprising and indicative of a likely change in the state of the environment or when their beliefs are more uncertain due to a lack of information about the current environmental state. Here we identified a pattern of whole-brain functional connectivity that changed dynamically over time, being more strongly expressed under conditions of high surprise or uncertainty, and being enhanced in individuals that adapted their learning and belief updating more appropriately in response to these factors. The key features of this whole-brain pattern of functional connectivity involved stronger connectivity, or functional integration, between the fronto-parietal and other functional systems. Our results provide new insights regarding the association between dynamic adjustments in learning and dynamic, large-scale changes in functional connectivity across the whole brain.

Introduction

Human decisions are guided by beliefs about the current environment. However, the underlying state of the environment is usually not directly observable and therefore people need to infer it from observable evidence that can be noisy. For example, deciding to go to a restaurant relies on your belief about the quality of the restaurant, which you infer based on past experiences at that restaurant. Furthermore, how much people update their beliefs in response to new evidence can depend on the uncertainty they have about their current beliefs and the extent to which the new evidence is surprising. If you have a lot of experience with a restaurant, and therefore less uncertainty in your beliefs about its quality, you may be more likely to discount a new experience that was better or worse than expected, compared to if you have little experience. Also, if a new experience is particularly surprising, it might indicate a fundamental change in the restaurant (e.g., a new chef) and subsequently spark a larger update in your belief. Past studies have shown that in dynamic environments people adaptively update their beliefs in a manner that depends on uncertainty and surprise^{1, 2, 3}. As both uncertainty and surprise increase, people update their beliefs in a way that weights current evidence more and discounts past experience to a greater degree. On the neural level, many studies have shown that such uncertainty and surprise in a dynamic environment was represented as univariate and multivariate activity in medial and lateral fronto-parietal networks^{1, 4, 5, 6}.

However, changes in belief updating may be reflected not only in changes in neural activity in single brain regions of fronto-parietal networks, but also in changes in functional connectivity between these brain regions. Functional connectivity reflects statistical dependencies between regional activity time series⁷ and provides a different perspective to understand the brain function^{8, 9, 10}. Recent studies have shown that network neuroscience approaches provide complementary measurements of the neural changes that occur during human learning¹¹. These studies focused on brain network reconfigurations that occur between naïve and well-learned phases in various domains such as motor, perceptual, category, spatial or value learning^{11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21}. In these cases, functional connectivity associated with fronto-parietal network decreased as learning progressed and this change in connectivity was associated with individual learning or performance^{11, 18, 21}. In dynamic environments, however, people progressively learn the current state and re-initialize learning once the state changes. Thus, we expect more rapid reconfigurations in functional connectivity, as learning shifts between slower and faster updating in response to changes in uncertainty and surprise.

In the current study, we aimed to identify such rapid reconfigurations in functional connectivity during adaptive belief updating. For this aim, we use an unsupervised machine learning technique known as non-negative matrix factorization (NMF)²². NMF decomposes the whole-brain network into subgraphs, which describe patterns of functional connectivity across the entire brain, and the time-dependent magnitude with which these subgraphs are expressed. Recently, NMF has been used to identify network dynamics during rest and task states^{23, 24}, and to determine how these dynamics vary across development²⁵. Here we extend the use of this technique to examine changes in network dynamics linked to task variables and individual differences. We hypothesized that uncertainty and surprise (which drive the adjustment of learning) are related to the temporal expression of specific patterns of functional connectivity (i.e., specific subgraphs) involving the fronto-parietal network. We also expected that the dynamic modulation of these patterns of functional connectivity (i.e., subgraph expression) are associated with individual differences in learning.

Results

Participants performed a predictive-inference task during fMRI (Fig. 1a). In this task, participants positioned a bucket to catch a bag that drops from an occluded helicopter. The location of the bag was sampled from a Gaussian distribution with a standard deviation (noise) and a mean (the location of the helicopter). The noise level was set to either high or low in each 120-trial run. The location of the helicopter usually remained stable but occasionally shifted suddenly and unpredictably (with an average probability of change of 0.1 across trials), with the new location sampled from a uniform distribution. Additionally, whether the bag (if caught) was rewarded or neutral was randomly assigned on each trial and indicated by color. This task challenged participants to form and update a belief about a hidden process (the location of the helicopter) based on noisy evidence (i.e., the location of dropped bags).

As described in our previous report¹, participants' predictions were influenced by both normative and non-normative factors. In a theoretical model approximating the Bayesian ideal observer, beliefs should be updated based on a delta-rule (Fig. 1b), with a learning rate (α_t) determined by two factors: change-point probability (CPP; i.e., the likelihood that a change-point has happened) and relative uncertainty (RU; i.e., the reducible uncertainty regarding the current state relative to the irreducible uncertainty due to environmental noise) (Fig. 1c). Namely, CPP and RU indicated belief surprise and belief uncertainty, respectively. Participants' learning increased as the values of these normative factors (CPP and RU) increased. Whether the outcome was rewarded, which is not a feature of the normative model, also influenced learning rates. We have previously reported how univariate and multivariate brain activity varies with CPP, RU and reward, and the relationship between univariate activations and individual differences^{1, 6}. In the current study, we investigated how the dynamics of whole-brain functional connectivity are related to these factors and to individual differences in belief updating.

Decomposing functional subgraphs

We used NMF to decompose whole-brain dynamic functional connectivity into specific patterns of functional connectivity, called subgraphs, and the expression of these patterns over time. We first identified regions of interest (ROIs) based on the parcellation in Power et al. (2011)²⁶ (Fig. 2a) and extracted blood-oxygenation-level-dependent (BOLD) time series for each ROI (Fig. 2b). For every ROI pair, we calculated the Pearson correlation coefficient between the BOLD time series in each 10-TR (25 seconds) time window. Consecutive time windows were offset from one another by 2 TRs, leading to 80% overlap between time windows. The correlation matrix in each time window (Fig. 2c) was then unfolded into a one-column vector. We concatenated these vectors from all time windows in all participants (Fig. 2d). The values of this matrix are required to be strictly positive for NMF, and, to ensure that our approach did not give undue preference to either positively or negatively weighted functional connectivity, we separated this matrix into two thresholded matrices. The full data matrix was divided into two halves, with one half containing the positive correlation coefficients (zero if the coefficient was negative) and one half containing the absolute values of negative correlation coefficients (zero if the coefficient was positive)²⁴. We applied NMF to this matrix to identify functional subgraphs and their expression over time. The full data matrix was decomposed into a matrix W , which reflected the strengths of edges for each subgraph, and a matrix H , which reflected the time-

dependent expression of each subgraph (Fig. 2d). The process of NMF was implemented through minimizing the residual error ($\|A - WH\|_F^2$) with the control of three parameters: the number of subgraphs (k), the subgraph regularization (α) and the expression sparsity (β) (Supplementary Fig. 1; see Methods). Through NMF, we reduced whole-brain dynamic functional connectivity into several subgraphs, which reflected patterns of functional connectivity over the whole brain, as well as the temporal expression of these subgraphs.

Properties of subgraphs

We summarized the properties of ten subgraphs we identified through NMF. Each subgraph was formed as edges between nodes (ROIs). We then converted these edges between nodes into edges between functional systems. We used 13 putative functional systems²⁶. Edges between nodes were categorized according to the corresponding system of each node. To estimate the diagonal entries in the system-by-system matrix, we averaged the weights of all edges within a given system (Fig. 3a). To estimate the off-diagonal entries of the system-by-system matrix, we averaged the weights of all edges linking a node in one system with a node in another system. In line with common parlance, we refer to the edges between two different systems as between-system edges whereas we refer to the edges within the same system as within-system edges. To demonstrate the connectivity pattern of each subgraph, we ordered them according to the strength of within-system edges relative to that of between-system edges (Fig. 3b, Supplementary Fig. 2a-c). For temporal expression, due to highly negative relationships between positive and negative expression (all $r < -0.61$, all $p < 0.001$), we used the relative expression (which is the difference between positive and negative expression) as the subgraph expression²⁴. Across subgraphs, the average expression was strongly correlated with relative strength of within- versus between-system edges (Supplementary Fig. 2d-f). That is, higher within-system strength was associated with greater temporal expression of the subgraph.

Modulation of temporal expression by factors that drive learning

We investigated the influence of learning factors on the temporal expression of each subgraph, and identified one subgraph in particular (subgraph 4) where expression was dynamically modulated by factors that should normatively drive learning. We focused on the two normative factors: CPP and RU. Theoretically, learning is dynamically driven by these two factors. We also included two other factors: reward and residual updating (i.e., belief updating not explained by the other factors). Through multiple regression, we estimated the influence of these four factors on trial-by-trial expression of each subgraph. Regression coefficients were fitted within each participant and each subgraph. Then, significance of these coefficients was tested at the group level for each subgraph (Supplementary Fig. 3). Among the ten subgraphs, these factors explained the most variance in the temporal expression of subgraph 4, and therefore we focused on this subgraph (Supplementary Fig. 2g). CPP (mean \pm s.e.m.=0.202 \pm 0.053; t_{31} =3.78, $p < 0.001$), RU (mean \pm s.e.m.=0.392 \pm 0.077; t_{31} =5.11, $p < 0.001$) and residual updating (mean \pm s.e.m.=0.177 \pm 0.079; t_{31} =2.23, $p=0.033$) all positively modulated the time-dependent expression of subgraph 4 (Fig 4a). Investigating the structure of this subgraph in more detail, its strongest edges involved the fronto-parietal task control system (Fig. 4d). No other subgraph showed such robust effects for both CPP and RU. These effects were unchanged when we

covaried for the influence of motion (indexed as the relative root-mean-square (RMS) of the six motion parameters) in the regression model.

The relationship between individual differences in learning and subgraph expression

The expression of subgraph 4 reflected individual difference in normative learning. First, we examined the relationship between normative learning and the dynamic modulation of subgraph expression by normative factors (Supplementary Fig. 4). For normative learning, we estimated the influence of CPP and RU on trial-by-trial belief updates through multiple regression and used the sum of the regression coefficients of CPP and RU as an index of normative learning for each participant¹ (see Methods). Similarly, we used the sum of regression coefficients of CPP and RU in the previous analyses (Supplementary Fig. 3) as the index of the dynamic modulation of subgraph expression by normative factors. For subgraph 4, there was a positive correlation between normative learning and dynamic modulation by normative factors across participants ($r=0.448$, $p=0.006$; Fig. 4b). Second, we examined the relationship between normative learning and the average subgraph expression (Supplementary Fig. 5). There was also a significantly positive correlation between normative learning and average expression of subgraph 4 across participants ($r=0.332$, $p=0.032$; Fig. 4c). These results show that participants with the highest average expression of subgraph 4, and for whom the normative factors account for the most variance in the expression of subgraph 4 across time, tended to update their beliefs in a manner more consistent with the normative model. Again, these effects were not changed when we covaried for the influence of motion (i.e., the relative RMS of the six motion parameters) in the regression model.

Contribution of specific edges to task and individual difference effects

Since subgraph 4 is characterized by balanced strength of within-system and between-system edges (Supplementary Fig. 2a-c), we evaluated the different contribution of within-system and between-system edges to the effects identified above. First, we determined the relative contribution of within-system and between-system edges to the dynamic modulation of subgraph expression by learning factors. We estimated the effects of CPP, RU, reward and residual updating on the expression of the full subgraph versus reduced subgraphs with only the within-system or between-system edges (Fig. 5a). Removing the between-system edges (Within versus All) reduced the size of the estimated effects of CPP (mean \pm s.e.m.=-0.155 \pm 0.042; $t_{31}=-3.73$, $p<0.001$), RU (mean \pm s.e.m.=-0.300 \pm 0.062; $t_{31}=-4.82$, $p<0.001$), and residual updating (mean \pm s.e.m.=-0.140 \pm 0.053; $t_{31}=-2.63$, $p=0.013$), while removing the within-system edges (Between versus All) led to no significant changes in these effects. Further, in a direct comparison of the reduced subgraphs with only within- or between-system edges, the effects estimated with between-system edges only were stronger for CPP (mean \pm s.e.m.=0.151 \pm 0.042; $t_{31}=3.63$, $p<0.001$), RU (mean \pm s.e.m.=0.290 \pm 0.063; $t_{31}=4.63$, $p<0.001$), and residual updating (mean \pm s.e.m.=0.139 \pm 0.048; $t_{31}=2.91$, $p=0.007$). Therefore, the effect of task factors on the temporal expression of subgraph 4 involved primarily between-systems connectivity.

The relationship between individual difference in normative learning and the dynamic modulation of subgraph 4 expression appeared to primarily involve within-system edges, while that with the average expression of subgraph 4 appeared to primarily involve between-system

edges, though neither of these conclusions was as statistically robust as that regarding task factors on subgraph expression (Fig. 5b). For the relationship between individual differences in normative learning and dynamic modulation, removing the within-system edges (Between versus All) slightly reduced the correlation (difference = 0.048, $p=0.006$). However, there was no significant difference for the direct comparison between between-system edges only and within-system edges only. On the other hand, for the relationship between individual differences in normative learning and average expression, there were no significant differences in any comparison, though the between-system edges seem to contribute more due to a weak reduction in the correlation after removing these edges (Within versus All) and a higher correlation when directly compared between-system edges only with within-system edges only.

We also investigated the contribution of the edges involving different functional systems or of specific system-by-system edges to the effects identified above. That is, we conducted similar analyses by removing all edges of one functional system (Supplementary Fig. 6) or removing one system-by-system edge (Supplementary Fig. 7). The strongest edges in subgraph 4 involve the fronto-parietal, memory retrieval, dorsal attention and salience systems (Fig. 4d). Removing edges associated with these systems led to significant reductions in the effects of CPP and RU on temporal expression of the subgraph (Supplementary Fig. 6-7). A similar, though weaker, relationship held for the contribution of specific systems/edges to the correlation between individual differences in normative learning and the dynamic modulation or average expression of the subgraph (Supplementary Fig. 6-7).

Robustness Checks

To determine the sensitivity of these results to the time window used to identify functional connectivity, we repeated the entire procedure after decomposing functional subgraphs using a shorter (8-TR/20 seconds window with 6-TR/15 seconds overlap; Supplementary Fig. 8-11) or longer time window (12-TR/30 seconds window with 10-TR/25 seconds overlap; Supplementary Fig. 12-15). In both shorter and longer time windows, we identified ten subgraphs. To delineate the similarities between these subgraphs and the identified subgraphs in the main analysis, we estimated the Pearson correlation coefficient between the edge strengths of these subgraphs and the edge strengths of their corresponding subgraphs in the main analysis. We repeated this analysis for edges between nodes and edges between systems. We found that these subgraphs were highly similar to those in the main analysis for both shorter (edges between nodes: all $r>0.81$; edges between systems: all $r>0.80$) and longer time windows (edges between nodes: all $r>0.98$; edges between systems: all $r>0.98$). The expression of subgraph 4 still showed the relationship to task factors (CPP and RU) and to individual differences in normative learning when identified with a longer time window; these effects were also present but weaker when identified with a shorter time window.

Relationship between univariate activation and functional connectivity

The nodes most strongly involved in subgraph 4 spatially overlapped regions where univariate activity was modulated by CPP and RU in our previous report. We calculated the edge strength of each ROI by averaging the interaction strength between this ROI with other ROIs in subgraph 4. Moreover, we extracted the z-statistic of the modulation effect of CPP and RU on

activation from our previous study¹. Across all ROIs, there was a positive correlation between edge strength and activation for CPP ($r=0.193$, $p<0.001$; Fig. 6a) and between edge strength and activation for RU ($r=0.394$, $p<0.0001$; Fig. 6b). We also compared the thresholded activation maps for CPP and RU with the map of normalized edge strength for subgraph 4 (Fig. 6c). Regions that show stronger increases in activation for increases in CPP and RU tended to also have high edge strength in subgraph 4, and *vice versa*.

Though there are strong associations between univariate activation and edge strength, our functional connectivity results are not entirely due to coordinated activations across regions that can be captured in univariate analyses. For each ROI, we constructed a time series of predicted BOLD activity from our univariate analyses, which reflected the fluctuations in activity in that ROI that can be accounted for by the four task factors (CPP, RU, reward and residual updating). Then, we implemented NMF on the matrix of functional connectivity estimated from the predicted BOLD time series and repeated all of our analyses (Supplementary Fig. 16-19). In the predicted BOLD, we again identified a subgraph 4 that showed the strongest edges in fronto-parietal system. However, only some of the effects identified above were present in the predicted BOLD. Namely, with regard to task effects, only RU significantly modulated the temporal expression of the subgraph, and regard to individual difference effects, only the relationship between individual differences in normative learning and average subgraph expression was significant. Thus, the predicted BOLD time series do not capture all the functional connectivity effects we identified in our main analyses. Instead, the dynamic functional connectivity patterns identified above seem to reflect a mixture of coordinated activity across regions (which can be captured by univariate analyses) and other statistical dependencies across regions which can only be captured with network analyses.

Discussion

We identified a pattern of functional connectivity that changed dynamically over time during a predictive inference task, being more strongly expressed during times that demanded faster learning and belief updating, and being enhanced in individuals that adapted their learning and belief updating more appropriately in this task. To identify this pattern, we used NMF, an unsupervised machine learning technique that decomposes the full matrix of time-dependent functional connectivity into subgraphs (patterns of functional connectivity), and the time-dependent magnitude of these subgraphs. Among the subgraphs we identified in our data, the expression of one subgraph in particular was coherently modulated by three trial-by-trial factors which behaviorally influenced the degree of belief updating: CPP (surprise), RU (uncertainty), and residual updating (updating unaccounted for by surprise or uncertainty). Notably, CPP and RU are factors that should normatively promote greater belief updating, discounting past observations relative to the most recent evidence. Residual updating reflects the deviation between the subject's update and the update that can be accounted for by the CPP and RU factors from the normative model. That is, the expression of this subgraph not only reflects normative factors but also likely fluctuations in subjective estimates of those factors. In addition, expression of this subgraph across individuals was statistically associated with individual differences in belief updating. Participants who showed stronger dynamic modulation of the expression of this subgraph by normative factors or who showed stronger average expression of this subgraph tended to update their beliefs in a more normative manner – that is, with a stronger influence of surprise (CPP) and uncertainty (RU).

Structure of the learning-related subgraph

We characterized the complex pattern of functional connectivity in the learning-related subgraph by summarizing the connectivity according to the putative functional system of each node²⁶. The largest proportion of connectivity in the learning-related subgraph involved the fronto-parietal system. Connectivity associated with the fronto-parietal system has been shown to increase at the beginning of learning and decreased toward the later phases of learning^{11, 18, 21}. Our result extends this finding and shows that fronto-parietal functional connectivity is dynamically modulated in a trial-by-trial manner according to the need for new learning. That is, the pattern of functional connectivity captured by the learning-related subgraph increased after state changes and then gradually decreased as more information was gained about the current state. The fronto-parietal system is thought of as a control system that is involved in adjusting behavior flexibly^{27, 28}. In particular, connectivity between the fronto-parietal network and other systems has been shown to change in response to different task requirements²⁹. This type of flexible control is critical for learning in a dynamic environment, a context in which people should adjust their degree of belief updating flexibly^{3, 4}.

The identified subgraph was also characterized by the balanced strength of within-system connectivity relative to between-system connectivity. However, we showed that the critical features that changed in response to task dynamics involved between-system connectivity. This result implies that faster learning and belief updating were associated with a greater degree of integration between different functional systems. Several previous studies have shown that complex cognitive tasks are associated with more integration between systems^{30, 31, 32, 33}. For example, brain networks tended to be more segregated during sensory processing, and more integrated during complex cognitive tasks such as the n-back or gambling. Other work has shown that as a task becomes more practiced over time that the interaction between systems decreased while the connections within systems remained strong¹¹. Here we demonstrated changes in integration on a fast time scale, as task demands varied from trial to trial. Integration between systems was greater during periods of the task when surprise or uncertainty was high, and therefore there was a need to update one's beliefs and base them more on the current evidence than on expectations developed from past experience.

Bridge between neural activation and neural communication

During learning in response to surprise and uncertainty, regions that showed stronger changes in neural activation also showed stronger functional connectivity. This association between neural activation and functional connectivity was also observed in other studies of motor, perceptual, category, spatial or value learning^{11, 14, 16, 18, 20, 21}. This pattern of results raises the possibility, though, that the regions involved in the subgraph may be becoming more tightly synchronized to external task events during periods of surprise and uncertainty, without necessarily any increase in communication between them.

However, this was not the case. The dynamic changes in functional connectivity we identified appeared to reflect a mixture of regional co-activation induced by task factors and increases in correlation between regions that could not be accounted for by task-related activations. The functional connectivity present in predicted BOLD time series from univariate

GLMs captured some, but not all, of the effects we observed in our data. There were both critical within-subject (e.g., the modulation of subgraph expression by surprise and residual updating) and between-subject (e.g., the relationship between individual differences in normative learning and the dynamic modulation of subgraph expression by normative factors) effects that were only uncovered in the full, original functional connectivity matrices. Even though the changes in functional connectivity we describe may reflect a mixture of stimulus-driven and endogenous dynamics, the network-level provides an important higher-level, reduced-dimensionality description of these changes.

Individual differences in learning were also reflected in features of individual connectomes. In our previous study, we noted a relationship between individual differences in normative learning and activity in dorsomedial frontal cortex and anterior insula¹. In the current study, we further identified the association between normative learning and a pattern of a whole-brain functional connectivity (both the average expression of this subgraph and the dynamic modulation of this subgraph's expression by normative factors). Thus, our present results show how functional connectivity can provide complimentary information to neural activations regarding individual differences. Our findings also add to previous work showing that brain network dynamics can reflect individual differences in learning in various domains^{11, 12, 14, 18, 21}. Potentially, these differences in individual connectomes during learning might be linked individual differences in resting-state functional connectivity (in which no task performance is required)³⁴.

Arousal and network dynamics

The dynamics of functional connectivity that we identified may be linked to arousal. We and others have previously shown that surprise and uncertainty induce changes in arousal, as indexed by pupil responses^{2, 5, 35}. The increases in integration between the fronto-parietal and other systems that we observed in this study in response to surprise and uncertainty are consistent with previous studies of the relationship between arousal (e.g., wake/sleep, heart rate and pupil size) and the dynamics of functional connectivity in large-scale brain networks^{33, 36, 37, 38, 39, 40}. Specifically, high arousal states are associated with changes in connectivity in regions in the fronto-parietal, dorsal attention, ventral attention, salience and default mode networks^{36, 39}, and particularly to increases in integration in the fronto-parietal and default mode networks³³.

Potential links between arousal and dynamic functional connectivity are particularly interesting in light of theories about the functional role of arousal in changing neural gain, or the sensitivity of neural communication⁴¹. Arousal is associated with changes in the activity of noradrenergic neurons in the locus coeruleus^{42, 43, 44}, which is hypothesized to modulate neural gain in the cortical projection targets of these cells. As neural gain increases, neurons receiving excitatory input become more active while neurons receiving inhibitory input become less active, and these effects have been shown to change network topology⁴⁵. One previous study established a simulated brain functional network and showed that increased neural gain leads to stronger network integration, which is particularly pronounced in the fronto-parietal network⁴⁶. Furthermore, atomoxetine, a noradrenergic reuptake inhibitor, increases network integration during task performance⁴⁷ and modulates participants' learning in dynamic environments⁴⁸.

Limitations and future directions

We conclude by noting a few limitations of the current study and several directions for future work. First, future studies should evaluate the generalizability of our results to other types of tasks. In the current study, we used a specific predictive-inference task in which the structure of the environment was communicated through visual representations. If instead the inference regarded a different modality such as auditory stimuli³⁵, how similar might the learning-related subgraph be? We would expect that different modalities require interactions between different sensory systems, but that the critical role of the fronto-parietal system would remain.

Second, future studies should investigate whether distinct network dynamics are associated with changes in surprise and uncertainty. In dynamic environments like the one we studied, the two factors are temporally linked, as change-points lead to immediate periods of surprise followed by longer periods of uncertainty¹. In the current study, functional connectivity was calculated in sliding time windows, which may include both the surprise elicited by change-points and the uncertainty that follows. Thus, surprise and uncertainty are difficult to temporally dissociate. This property may explain why we identified a single subgraph associated with both factors. Our previous work showed that surprise was selectively associated with activation in visual cortex and uncertainty was selectively associated with activation in anterior prefrontal and parietal cortex¹. Using a task that can temporally separate the tracking of surprise and uncertainty⁴⁹ might enable the identification of distinct subgraphs for each factor.

Last, future studies should evaluate the distinct influence of arousal on the dynamics of functional connectivity during belief updating. Arousal may play an important role during periods of high cognitive demand when there is a need for the flexible adjustment of behavior. There were important similarities between the learning-related subgraph we identified and previously described arousal-linked changes in functional integration. However, we did not directly measure arousal in the current study. It would be important to investigate the relationship between arousal, network reconfiguration, and belief updating in the future.

Methods

Participants

We used a dataset from our previous study and further details can be found in that report¹. Thirty-two individuals participated in the fMRI experiment: 17 females, mean age = 22.4 years (SD = 3.0; range 18-30). Human subject protocols were approved by the Internal Review Board in University of Pennsylvania. All participants provided informed consent before the experiment.

Task

Each participant completed four 120-trial runs during functional scanning. In each run, participants performed a predictive-inference task (Fig. 1a). On each trial, participants made a prediction about where the next bag would be dropped from an occluded helicopter by positioning a bucket along the horizontal axis (0-300) of the screen. The location of the bag was sampled from a Gaussian distribution with a standard deviation (noise) and a mean (the location of the helicopter). The standard deviation was high (s.d. = 25) or low (s.d. = 10) in different runs. The location of the helicopter usually remained stable but it changed occasionally. The probability of change was zero for the first three trials after a change and 0.125 for the following trials. After the change, the new location was sampled from a uniform distribution. Additionally, correctly predicting the location of the bag resulted in coins landing in the bucket, which could either have positive or neutral value depending on their color, which was randomly assigned for each trial.

Behavior model

We applied the same normative model described in our previous study¹. An approximation to the ideal observer solution to this task updates beliefs according to a delta learning rule (Fig. 1b). Beliefs are updated in proportion to the prediction error, which is the difference between the observed outcome (bag drop location) and the prediction (bucket location), $X_t - B_t$. The proportion of the prediction error to update is determined by the learning rate (α_t), which is adjusted adaptively on each trial according to two normative factors (Fig. 1c): change-point probability (CPP) and relative uncertainty (RU). CPP reflects the likelihood that a change-point has happened and increases when there is an unexpectedly large prediction error. RU reflects the uncertainty about the current location of the helicopter relative to the amount of noise in the environment. RU increases after CPP increases and decays slowly as more precise estimates of the helicopter location are possible.

As in our previous study, a regression model was applied to investigate how different factors influence participants' belief updates. We regressed trial-by-trial updates ($B_{t+1} - B_t$) against the prediction error (δ) and its interaction with the two normative factors, CPP and RU, as well as whether the outcome was rewarded or not¹. The form of the regression model can be written as

$$Update_t = \beta_0 + \beta_1 \delta_t + \beta_2 \delta_t CPP_t + \beta_3 \delta_t RU_t (1 - CPP_t) + \beta_4 \delta_t Reward_t + \beta_5 Edge_t + \varepsilon ,$$

where t is the trial number and $Edge$ is a quadratic-weighted term $((150 - B_{t+1})|150 - B_{t+1}|)$ that captures the tendency to avoid updating toward the edges of the screen. This regression model was fitted separately to each participant's data to estimate the influence of each factor on each participant's behavior. To examine the relationship between individual differences in normative learning and functional network dynamics, we used the sum of the regression coefficients on CPP and RU as an index of normative learning.

MRI data acquisition and preprocessing

MRI data were collected on a 3T Siemens Trio with a 32-channel head coil. Functional data were acquired using a gradient-echo echoplanar imaging (EPI) (3 mm isotropic voxels, 64 x 64 matrix, 42 axial slices tilted 30° from the AC-PC plane, TE = 25 ms, flip angle = 75°, TR = 2500 ms). There were four runs with 226 images per run. T1-weighted MPRAGE structural images (0.9375 X 0.9375 X 1 mm voxels, 192 X 256 matrix, 160 axial slices, TI = 1100 ms, TE = 3.11 ms, flip angle = 15°, TR = 1630 ms) and matched fieldmap images (TE = 2.69 and 5.27 ms, flip angle = 60°, TR = 1000 ms) were also collected. Functional data were corrected for slice timing and head motion, attenuated for outliers, undistorted and warped to MNI space, smoothed with 6 mm FWHM Gaussian kernel and intensity-scaled by the grand-mean value per run. Detail of these preprocessing steps were described in our previous study¹. Individual's structure images were segmented into gray matter, white matter (WM) and cerebrospinal fluid (CSF) using FAST (FMRIB's Automated Segmentation Tool)⁵⁰.

Constructing time-varying functional networks

For each run and each participant, blood-oxygenation-level-dependent (BOLD) time series were obtained from each of 264 regions of interest (ROIs; diameter = 9mm) based on the parcellation in Power et al. (2011)²⁶. ROIs that did not have valid BOLD time series for all runs and all participants were removed, resulting in $N = 247$ ROIs. For each BOLD time series, a band-pass filter was applied with a cutoff of 0.01-0.08 Hz. This low-frequency band has shown to reflect neuronal activation and neural synchronization^{51, 52, 53}. To remove the influence of head motion, a confound regression was implemented to regress out nuisance factors from each BOLD time series. This confound regression included 24 motion parameters (three translation and three rotation motion parameters and their expansion $([R_t R_t^2 R_{t-1} R_{t-1}^2])$)⁵⁴, as well as average signals from WM and CSF⁵⁵.

In order to construct dynamic functional networks, we define sliding time windows and calculate Pearson correlation coefficients between ROI time series in each sliding time window. We assigned these coefficients to the first TR in the time windows. To ensure magnetization equilibrium, the first 6 volumes of each run were removed from the analysis. For the rest of the volumes in each run, a sliding window was defined with a 10-TR (i.e., 25 seconds) length and 80% overlap across windows. Each run had 106 sliding time windows, leading to $T = 424$ sliding time windows for each participant. Each participant's data thus formed a matrix of dynamic functional networks with dimensions $N \times N \times T$. Then, we took each participant's $N \times N$ matrix and unfurled the upper triangle into an $\frac{N(N-1)}{2}$ vector. By concatenating vectors across all time windows (T), we obtained an $\frac{N(N-1)}{2} \times T$ matrix. Furthermore, we concatenated matrices from $S =$

32 participants to form a $\frac{N(N-1)}{2} \times (T \times S)$ matrix. To ensure that our approach did not give undue preference to either positively or negatively weighted functional edges, we separated this matrix into two thresholded matrices: one composed of positively weighted edges, and one composed of negatively weighted edges. That is, in the matrix of positive functional correlations between ROI time series, the original negative correlations between ROI time series were set to 0; in the matrix of negative functional correlations between ROI time series, all values were multiplied by -1, and the original positive functional correlations between ROI time series were set to 0. After concatenating the matrix composed of positively weighted edges and the matrix of negatively weighted edges, we had a final $\frac{N(N-1)}{2} \times (T \times S \times 2)$ matrix \mathbf{A} .

Clustering functional networks into subgraphs

An unsupervised machine learning algorithm called non-negative matrix factorization (NMF)²² was applied on \mathbf{A} to identify subgraphs \mathbf{W} and the time-dependent expressions of subgraphs \mathbf{H} . The matrix factorization problem $\mathbf{A} \approx \mathbf{WH}$ s. t. $\mathbf{W} \geq 0, \mathbf{H} \geq 0$ was solved by optimization of the cost function:

$$\min_{\mathbf{W}, \mathbf{H}} \frac{1}{2} \|\mathbf{A} - \mathbf{WH}\|_F^2 + \alpha \|\mathbf{W}\|_F^2 + \beta \sum_{t=1}^{TS} \|\mathbf{H}(:, t)\|_1^2,$$

where \mathbf{A} is the functional connectivity matrix, \mathbf{W} is a matrix of subgraph connectivity with size $\frac{N(N-1)}{2} \times k$, and \mathbf{H} is a matrix of time-dependent expression coefficients for subgraphs with size $k \times (T \times S)$. The parameter k is the number of subgraphs, α is a regularization of the connectivity for subgraphs, and β is a penalty that imposes sparsity on the temporal expression coefficients⁵⁶. We used an alternative non-negative least square with block-pivoting method with 100 iterations for fast and efficient factorization to solve this equation⁵⁷. The matrices \mathbf{W} and \mathbf{H} were initialized with randomized values from a uniform distribution between 0 and 1.

A random sampling procedure was used to find the optimal parameters k , α , and β ⁵⁸. In this procedure, the NMF algorithm was re-run 1,000 times with parameter k drawn from $U(2, 15)$, parameter α drawn from $U(0.01, 1)$, and parameter β drawn from $U(0.01, 1)$. The subgraph learning performance was evaluated through 4-fold cross-validation. In each fold, twenty-four participants were used for training; 8 participants were used for testing and calculating cross-validation error ($\|\mathbf{A} - \mathbf{WH}\|_F^2$). An optimal parameter set should minimize the cross-validation error. We chose an optimal parameter set ($k = 10, \alpha = 0.535, \beta = 0.230$) that ensured the cross-validation error in the bottom 25% of the distribution of cross-validation error from our random sampling scheme²³.

Since the result of NMF is non-deterministic, we implemented consensus clustering to obtain reliable subgraphs⁵⁹. In this procedure, we (i) used the optimal parameters and ran the NMF 100 times on \mathbf{A} , (ii) concatenated subgraph matrix \mathbf{W} across 100 runs into an aggregate matrix with dimensions $\frac{N(N-1)}{2} \times (k \times 100)$, (iii) applied NMF to this aggregate matrix to obtain a final set of subgraphs $\mathbf{W}_{consensus}$ and expression coefficients $\mathbf{H}_{consensus}$.

Properties of subgraphs

After applying NMF, we obtained the pattern of functional connectivity (W) and temporal expression (H) of subgraphs. To delineate interactions between nodes, we rearrange W into several $N \times N$ subgraphs. In addition, to understand the roles of cognitive systems in each subgraph, we mapped each node to 13 putative cognitive systems: uncertain, sensory, cingulo-opercular task control, auditory, default mode, memory retrieval, visual, fronto-parietal task control, salience, subcortical, dorsal attention, ventral attention, and cerebellar^{25, 26}. To show which within-system and between-system edges were strong in each subgraph, we applied a permutation test. We permuted the system label for nodes and formed a matrix with system-by-system edges. This process was repeated 10,000 times to determine which strength of system-by-system edges was above the 95% confidence interval threshold after correction for multiple comparisons.

To demonstrate the connectivity pattern of each subgraphs, we ordered them according to the relative strength between within-system edges and between-system edges. For each subgraph, we calculated the average strength of within-system edges (edges that link two nodes that both belong to the same system), and the average strength of between-system edges (edges that link a node in one system to a node in another system). Then, we subtracted the average strength of between-system edges (E_B) from the average strength of within-system edges (E_W) and divided this difference by the sum of them ($\frac{E_W - E_B}{E_W + E_B}$). We then summarize the properties of each subgraph by their relative strength, average within-system strength and average between-system strength. Additionally, we estimated their 95% confidence interval by implementing bootstrapping 10,000 times on the edges corresponding to the relative strength, average within-system strength or average between-system strength, separately.

Next, we investigated the relationship between the connectivity pattern and temporal expression of subgraphs. Due to high associations between positive and negative expression, we used the relative expression (which is the difference between positive and negative expression) as the subgraph expression²⁴. Across subgraphs, we then calculated Pearson correlation coefficients between the average expression and the relative strength, between the average expression and the average within-system strength, and between the average expression and the average between-system strength. To determine the significance of correlation coefficients, we implemented permutation tests. For each analysis, we permuted the labels of subgraphs for the average expression and calculated the correlation coefficient, and we repeated this procedure 10,000 times to form the null distribution of correlation coefficients. Then, we calculated p values by comparing the true correlation coefficient with the null distribution.

Modulation of learning factors on temporal expression

We investigated how fluctuations in the trial-by-trial expression of each subgraph were related to trial-by-trial learning factors. Through NMF, we acquired temporal expression every 2 TR (i.e., 5 seconds). We then applied a linear interpolation on the temporal expression to obtain an expression aligned with outcome onset on each trial. That is, we examined how different learning factors lead to the change of functional connectivity after the onset of outcome on each trial. Several learning factors were examined: CPP, RU, reward, and residual updating. The factors CPP and RU were estimated based on the normative learning model^{1, 2, 3}. Residual updating was derived as the residual of the behavioral regression model, which was the difference between a participant's actual update and the update predicted by the behavioral

regression model. We next examined the effect of these four factors together, by including all four factors in a regression model predicting trial-by-trial expression. Regression models were implemented for each participant separately. Regression coefficients were then tested at the group level using two-tailed *t*-tests.

The relationship between individual normative learning and subgraph expression

Next, we examined the relationship between individual differences in normative learning and subgraph expression. For the index of individual normative learning, we used the sum of the regression coefficients of CPP and RU estimated from the behavior model as an index of normative learning¹. This normative learning index reflected the extent to which a participant's trial-by-trial updates were influenced by the two normative factors CPP and RU. Then, we examined two types of relationship. First, we calculated the Pearson correlation coefficient between individual normative learning and dynamic modulation of subgraph expression by normative factors. This dynamic modulation was indexed as the sum of regression coefficients of CPP and RU estimated from the regression model for trial-by-trial subgraph expression. That is, this dynamic modulation reflected how normative factors were associated with the change in subgraph expression. Second, we calculated the Pearson correlation coefficient between individual normative learning and the average expression of subgraphs. To determine the significance of correlation coefficients, we implemented permutation tests. For each analysis, we permuted the labels of participants for individual normative learning and calculated the correlation coefficient, and we repeated this procedure 10,000 times to form the null distribution of correlation coefficients. Then, we calculated *p* values by comparing the true correlation coefficient with the null distribution.

Within-subject versus between-subject effects

To show the relative importance of each subgraph, we evaluated the variance explained for within-subject and between-subject effects for each subgraph. The explained variance of within-subject effects was the R^2 of the regression model that investigated the relationship between the four learning factors (CPP, RU, reward and residual updating) and trial-by-trial subgraph expression. The explained variance of between-subject effects was the R^2 of a regression model that investigated the relationship between individual differences in normative learning and two individual-level measures of subgraph expression (dynamic modulation by normative learning factors and average expression). We investigated the relationship between within-subject and between-subject effects across subgraphs using Pearson correlation. To determine the significance of the correlation coefficient, we implemented a permutation test. We permuted the labels of subgraphs for the within-subject effects and calculated its correlation coefficient with between-subject effects, and we repeated this procedure 10,000 times to form the null distribution of correlation coefficients. Then, we calculated *p* values by comparing the true correlation coefficient with the null distribution.

Contributions of different types of edges to within-subject and between-subject effects

We evaluated the contributions of different types of edges to within-subject and between-subject effects. We mainly focused on the contribution of within-system edges and between-

system edges. For this analysis, we implemented three types of comparison: Within versus All, Between versus All, and Between versus Within. For Within versus All, we kept within-system edges only and re-estimated within-subject and between-subject effects; then, we compared these effects with the effects estimated using all edges. This comparison showed the change of effects after between-system edges were removed, and thus, this comparison revealed the contribution of between-system edges. Specifically, for the within-subject effects, we examined the change of coefficients for the regression model that investigate the influence of four learning factors (CPP, RU, reward and residual updating) on subgraph expression. The change was calculated for each participant separately, and the significance of change was then tested at the group level using two-tailed *t*-tests. For the between-subject effects, we examined the change of correlation coefficients for two types of relationship separately: the relationship between individual normative learning and dynamic modulation and the relationship between individual normative learning and average expression. To determine the significance of the change of correlation coefficients, we implemented permutation tests. We permuted the labels of participants for individual normative learning and calculated the change of correlation coefficient, and we repeated this procedure 10,000 times to form the null distribution of the change of correlation coefficients. Then, we calculated *p* values by comparing the true change of correlation coefficient with the null distribution. Similar analyses were implemented for Between versus All. We kept between-system edges only and re-estimated within-subject and between-subject effects. We then compared these effects with the effects estimated using all edges. In this comparison, within-system edges were removed and thus, we examined the contribution of within-system edges. Last, the comparison of Between versus Within is a direction comparison between effects estimated with between-system edges only and effects estimated with within-system edges only. Thus, this comparison examined the different contributions between between-system and within-system edges.

We also investigated the contribution of different functional systems and the contribution of different system-by-system edges. For the contribution of different functional systems, we compared the effects after removing edges of one functional system with the effects estimated with all edges. For the contribution of different system-by-system edges, we compared the effects after removing one system-by-system edge with the effects estimated with all edges. Statistical testing was conducted with the same procedures.

Relationship between univariate activation and functional connectivity

We investigated the relationship between univariate activation and functional connectivity. To address this issue, we first estimated the modulation of univariate activation by four learning factors (CPP, RU, reward and residual updating) using general linear models (GLMs). In the GLM, the regressors were the onset of outcome, CPP, RU, reward and residual updating. Before entering the GLM, these regressors were convolved with a gamma hemodynamic response function as well as the temporal derivative of this function. Six motion parameters were also included in the GLM. After implementing the GLMs on univariate activity, we calculated the predicted BOLDs as the sum of predicted time series (including both the main effect and the temporal derivative) of CPP, RU, reward and residual updating. Then, we repeated the same sequence of analyses described above on the predicted BOLD time series.

Data availability

The data for the current study are available from the corresponding author upon request.

Code availability

Codes are available at https://github.com/changhaokao/nmf_network_learning.

References

1. McGuire JT, Nassar MR, Gold JI, Kable JW. Functionally Dissociable Influences on Learning Rate in a Dynamic Environment. *Neuron* **84**, 870-881 (2014).
2. Nassar MR, Rumsey KM, Wilson RC, Parikh K, Heasley B, Gold JI. Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience* **15**, 1040-1046 (2012).
3. Nassar MR, Wilson RC, Heasley B, Gold JI. An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *The Journal of Neuroscience* **30**, 12366-12378 (2010).
4. Behrens TEJ, Woolrich MW, Walton ME, Rushworth MFS. Learning the value of information in an uncertain world. *Nat Neurosci* **10**, 1214-1221 (2007).
5. O'Reilly JX, Schüffelgen U, Cuell SF, Behrens TEJ, Mars RB, Rushworth MFS. Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences* **110**, E3660-E3669 (2013).
6. Nassar MR, McGuire JT, Ritz H, Kable JW. Dissociable Forms of Uncertainty-Driven Representational Change Across the Human Brain. *The Journal of Neuroscience* **39**, 1688-1698 (2019).
7. Friston KJ. Functional and Effective Connectivity: A Review. *Brain Connectivity* **1**, 13-36 (2011).
8. Bassett DS, Sporns O. Network neuroscience. *Nat Neurosci* **20**, 353-364 (2017).
9. Bullmore E, Sporns O. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience* **10**, 186 (2009).
10. Medaglia JD, Lynall M-E, Bassett DS. Cognitive network neuroscience. *Journal of cognitive neuroscience* **27**, 1471-1491 (2015).
11. Bassett DS, Yang M, Wymbs NF, Grafton ST. Learning-induced autonomy of sensorimotor systems. *Nat Neurosci* **18**, 744-751 (2015).
12. Bassett DS, Wymbs NF, Porter MA, Mucha PJ, Carlson JM, Grafton ST. Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences* **108**, 7641-7646 (2011).
13. Mattar MG, Wymbs NF, Bock AS, Aguirre GK, Grafton ST, Bassett DS. Predicting future learning from baseline network architecture. *NeuroImage* **172**, 107-117 (2018).

14. Gerraty RT, Davidow JY, Foerde K, Galvan A, Bassett DS, Shohamy D. Dynamic Flexibility in Striatum-Cortical Circuits Supports Reinforcement Learning. *The Journal of Neuroscience* **38**, 2442-2453 (2018).
15. Mattar MG, Thompson-Schill SL, Bassett DS. The network architecture of value learning. *Network Neuroscience* **2**, 128-149 (2018).
16. Antzoulatos Evan G, Miller Earl K. Increases in Functional Connectivity between Prefrontal Cortex and Striatum during Category Learning. *Neuron* **83**, 216-225 (2014).
17. Baeg EH, Kim YB, Kim J, Ghim J-W, Kim JJ, Jung MW. Learning-Induced Enduring Changes in Functional Connectivity among Prefrontal Cortical Neurons. *The Journal of Neuroscience* **27**, 909-918 (2007).
18. Büchel C, Coull JT, Friston KJ. The predictive value of changes in effective connectivity for human learning. *Science* **283**, 1538-1541 (1999).
19. Fatima Z, Kovacevic N, Masic B, McIntosh AR. Dynamic functional connectivity shapes individual differences in associative learning. *Human Brain Mapping* **37**, 3911-3928 (2016).
20. Lewis CM, Baldassarre A, Committeri G, Romani GL, Corbetta M. Learning sculpts the spontaneous activity of the resting human brain. *Proceedings of the National Academy of Sciences* **106**, 17558-17563 (2009).
21. Sun FT, Miller LM, Rao AA, D'esposito M. Functional connectivity of cortical networks involved in bimanual motor sequence learning. *Cerebral cortex* **17**, 1227-1234 (2006).
22. Lee DD, Seung HS. Learning the parts of objects by non-negative matrix factorization. *Nature* **401**, 788-791 (1999).
23. Khambhati AN, Mattar MG, Wymbs NF, Grafton ST, Bassett DS. Beyond modularity: Fine-scale mechanisms and rules for brain network reconfiguration. *NeuroImage* **166**, 385-399 (2018).
24. Khambhati AN, Medaglia JD, Karuza EA, Thompson-Schill SL, Bassett DS. Subgraphs of functional brain networks identify dynamical constraints of cognitive control. *PLOS Computational Biology* **14**, e1006234 (2018).
25. Chai LR, *et al.* Evolution of brain network dynamics in neurodevelopment. *Network Neuroscience* **1**, 14-30 (2017).
26. Power JD, *et al.* Functional Network Organization of the Human Brain. *Neuron* **72**, 665-678 (2011).

27. Cole MW, Repovš G, Anticevic A. The Frontoparietal Control System: A Central Role in Mental Health. *The Neuroscientist* **20**, 652-664 (2014).
28. Power JD, Petersen SE. Control-related systems in the human brain. *Current Opinion in Neurobiology* **23**, 223-228 (2013).
29. Cole MW, Reynolds JR, Power JD, Repovš G, Anticevic A, Braver TS. Multi-task connectivity reveals flexible hubs for adaptive task control. *Nature neuroscience* **16**, 1348-1355 (2013).
30. Bertolero MA, Yeo BTT, D'Esposito M. The modular and integrative functional architecture of the human brain. *Proceedings of the National Academy of Sciences* **112**, E6798-E6807 (2015).
31. Shine JM, Poldrack RA. Principles of dynamic network reconfiguration across diverse brain states. *NeuroImage*, (2017).
32. Cohen JR, D'Esposito M. The Segregation and Integration of Distinct Brain Networks and Their Relationship to Cognition. *The Journal of Neuroscience* **36**, 12083-12094 (2016).
33. Shine JM, *et al.* The Dynamics of Functional Brain Networks: Integrated Network States during Cognitive Task Performance. *Neuron* **92**, 544-554 (2016).
34. Tompson S, Falk EB, Vettel JM, Bassett DS. Network Approaches to Understand Individual Differences in Brain Connectivity: Opportunities for Personality Neuroscience. *Personality Neuroscience*, (2018).
35. Krishnamurthy K, Nassar MR, Sarode S, Gold JJ. Arousal-related adjustments of perceptual biases optimize perception in dynamic environments. *Nature Human Behaviour* **1**, 0107 (2017).
36. Wang C, Ong JL, Patanaik A, Zhou J, Chee MWL. Spontaneous eyelid closures link vigilance fluctuation with fMRI dynamic connectivity states. *Proceedings of the National Academy of Sciences* **113**, 9653-9658 (2016).
37. Chang C, *et al.* Tracking brain arousal fluctuations with fMRI. *Proceedings of the National Academy of Sciences*, (2016).
38. Haimovici A, Tagliazucchi E, Balenzuela P, Laufs H. On wakefulness fluctuations as a source of BOLD functional connectivity dynamics. *Scientific Reports* **7**, 5908 (2017).
39. Patanaik A, Tandji J, Ong JL, Wang C, Zhou J, Chee MWL. Dynamic functional connectivity and its behavioral correlates beyond vigilance. *NeuroImage* **177**, 1-10 (2018).

40. Young CB, *et al.* Dynamic Shifts in Large-Scale Brain Network Balance As a Function of Arousal. *The Journal of Neuroscience* **37**, 281-290 (2017).
41. Servan-Schreiber D, Printz H, Cohen JD. A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior. *Science* **249**, 892-895 (1990).
42. Aston-Jones G, Cohen JD. An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annual Review of Neuroscience* **28**, 403-450 (2005).
43. Sara Susan J. The locus coeruleus and noradrenergic modulation of cognition. *Nat Rev Neurosci* **10**, 211-223 (2009).
44. Sara Susan J, Bouret S. Orienting and Reorienting: The Locus Coeruleus Mediates Cognition through Arousal. *Neuron* **76**, 130-141 (2012).
45. Eldar E, Cohen JD, Niv Y. The effects of neural gain on attention and learning. *Nat Neurosci* **16**, 1146-1153 (2013).
46. Shine JM, Aburn MJ, Breakspear M, Poldrack RA. The modulation of neural gain facilitates a transition between functional segregation and integration in the brain. *Elife* **7**, e31130 (2018).
47. Shine JM, van den Brink RL, Hernaus D, Nieuwenhuis S, Poldrack RA. Catecholaminergic manipulation alters dynamic network topology across cognitive states. *Network Neuroscience* **2**, 381-396 (2018).
48. Jepma M, Murphy PR, Nassar MR, Rangel-Gomez M, Meeter M, Nieuwenhuis S. Catecholaminergic Regulation of Learning Rate in a Dynamic Environment. *PLOS Computational Biology* **12**, e1005171 (2016).
49. Nassar MR, Bruckner R, Gold JI, Li S-C, Heekeren HR, Eppinger B. Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature communications* **7**, 11609 (2016).
50. Zhang Y, Brady M, Smith S. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging* **20**, 45-57 (2001).
51. Biswal B, Zerrin Yetkin F, Haughton VM, Hyde JS. Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine* **34**, 537-541 (1995).
52. Lu H, *et al.* Synchronized delta oscillations correlate with the resting-state functional MRI signal. *Proceedings of the National Academy of Sciences* **104**, 18265-18269 (2007).

53. Zuo X-N, *et al.* The oscillating brain: Complex and reliable. *NeuroImage* **49**, 1432-1445 (2010).
54. Friston KJ, Williams S, Howard R, Frackowiak RSJ, Turner R. Movement-Related effects in fMRI time-series. *Magnetic Resonance in Medicine* **35**, 346-355 (1996).
55. Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, Raichle ME. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the United States of America* **102**, 9673-9678 (2005).
56. Kim J, Park H. Fast Nonnegative Matrix Factorization: An Active-Set-Like Method and Comparisons. *SIAM Journal on Scientific Computing* **33**, 3261-3281 (2011).
57. Kim J, He Y, Park H. Algorithms for nonnegative matrix and tensor factorizations: a unified view based on block coordinate descent framework. *Journal of Global Optimization* **58**, 285-319 (2014).
58. Bergstra J, Bengio Y. Random search for hyper-parameter optimization. *Journal of Machine Learning Research* **13**, 281-305 (2012).
59. Monti S, Tamayo P, Mesirov J, Golub T. Consensus Clustering: A Resampling-Based Method for Class Discovery and Visualization of Gene Expression Microarray Data. *Machine Learning* **52**, 91-118 (2003).

Acknowledgements

This work was supported by grants from National Institute of Mental Health (R01-MH098899 to J.I.G. and J.W.K.) and National Science Foundation (1533623 to J.I.G. and J.W.K.). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

C.-H.K. and J.W.K. designed the study. A.N.K. established the software for NMF. A.N.K. and D.S.B. provided suggestions in network analyses. C.-H.K. implemented all the analyses and visualization and drafted the manuscript. All authors interpreted the results and revised the manuscript.

Competing interests

The authors declare no competing interests.

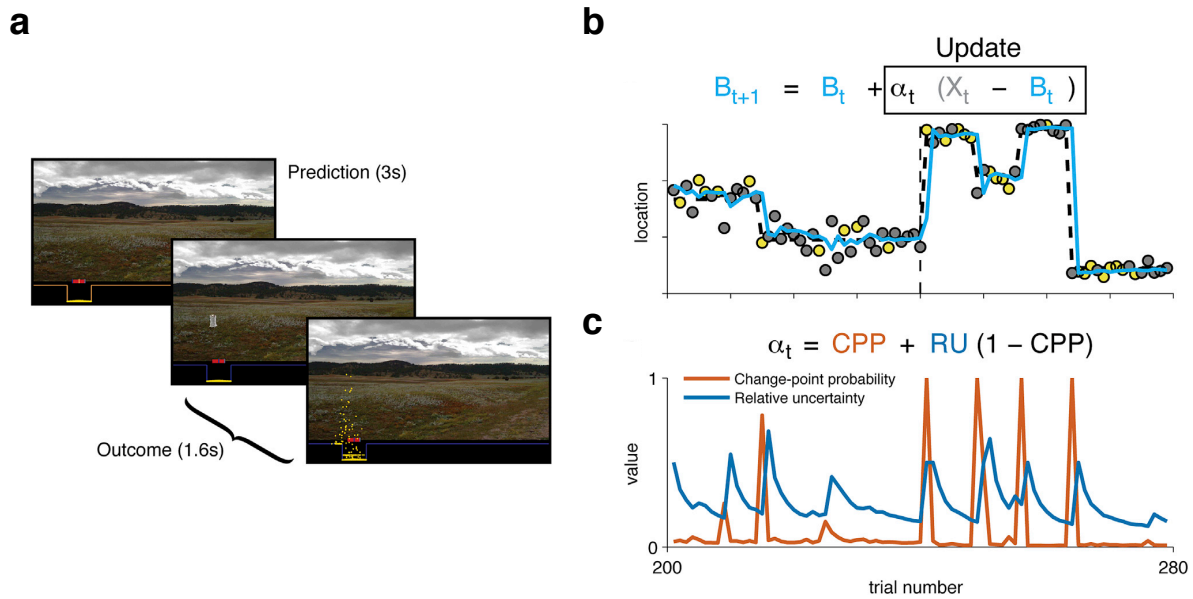


Figure 1. Overview of the task and theoretical model of belief updating (McGuire et al., 2014). (a) Sequence of the task. At the start of each trial, participants predict where a bag will drop from an occluded helicopter by positioning a bucket on the screen. After participants submit their prediction, the bag drops and any rewarded coins that fall in the bucket are added to the participant's score. The location of the last prediction and the last bag drop are noted on the next trial. (b) An example sequence of trials. Each data point represents the location of a bag on each trial (yellow for rewarded coins, gray for neutral coins). The dashed line represents the true generative mean. The mean changes occasionally. The cyan line represents the prediction from a normative model of belief updating. The inset equation shows how the model updates beliefs (B_t = belief, X_t = observed outcome, α_t = learning rate on trial t). The vertical dashed line represents the boundary of the noise conditions: high-noise (left) and low-noise condition (right). Noise refers to the variance of the generative distribution. (c) Two learning components from the normative model. Change-point probability (CPP) reflects the likelihood that a change-point happens, which is increased when there is an unexpectedly large prediction error. Relative uncertainty (RU) reflects the uncertainty about the generative mean relative to the environmental noise, which is increased after high CPP trials and decays slowly as more precise estimates of the generative mean are possible. The inset formula shows how CPP and RU contribute to single trial estimates of learning rates.

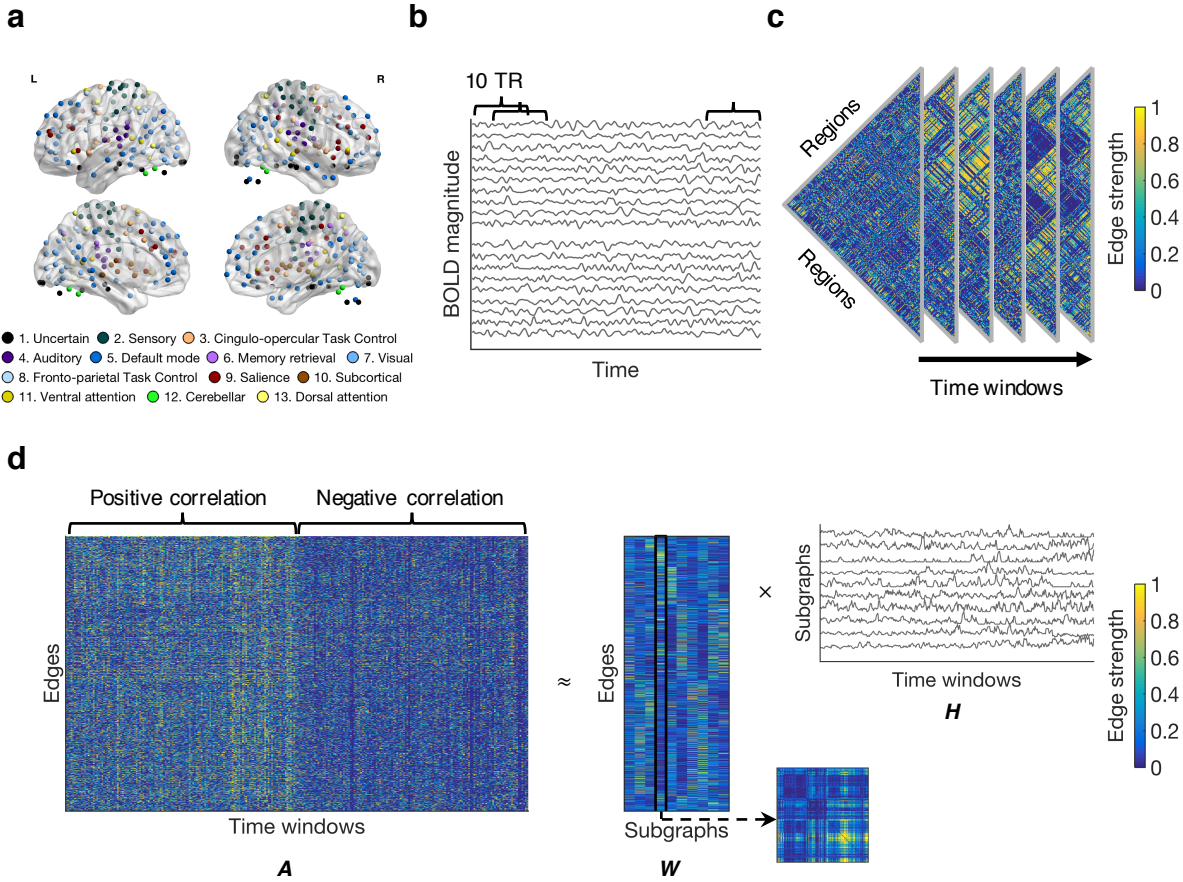


Figure 2. Schematic overview of the method. (a) Regions of interest (ROIs). Functional MRI BOLD signals were extracted from spherical ROIs based on the coordinate system of Power et al. (2011). There were 264 ROIs in this coordinate system. We only kept ROIs that had usable data from all subjects. Thus, we used 247 ROIs for the following analyses. Each ROI can be assigned to one of 13 putative functional systems. (b) An example of Pearson correlation coefficients calculated between regional BOLD time series over the course of the experiment. Each BOLD time series was divided into 10-TR (25 seconds) time windows, and consecutive time windows were placed every 2 TRs leading to 80% overlap between consecutive time windows. Pairwise Pearson correlation coefficients were calculated between ROI time series in each time window. (c) An example of edge strength over time. In each time window, there were $247 \cdot (247 - 1) / 2$ edges. (d) Nonnegative matrix factorization (NMF). In each time window, the matrix of edge strengths was unfolded into one column. Then, edges from all time windows in all participants were concatenated into a single matrix. Each row in the full data matrix contained an edge (pairwise correlation coefficients between BOLD time series from two ROIs) and each column contained a time window (across all scans and participants). Correlation values in this matrix were strictly positive; the full data matrix was divided into two halves, with one half containing the positive pairwise correlation coefficients (zero if the correlation coefficient was negative) and one half containing the absolute values of negative pairwise correlation coefficients (zero if the correlation coefficient was positive). Then, NMF was applied to decompose the concatenated matrix into a matrix W , which encoded the strengths of edges for each subgraph, and a matrix H , which encoded the time-dependent expression of each subgraph. For example, the strength of edges of the fourth subgraph (the fourth column in the matrix W) can be folded into a squared matrix, reflecting the edge strength between every pair of ROIs.

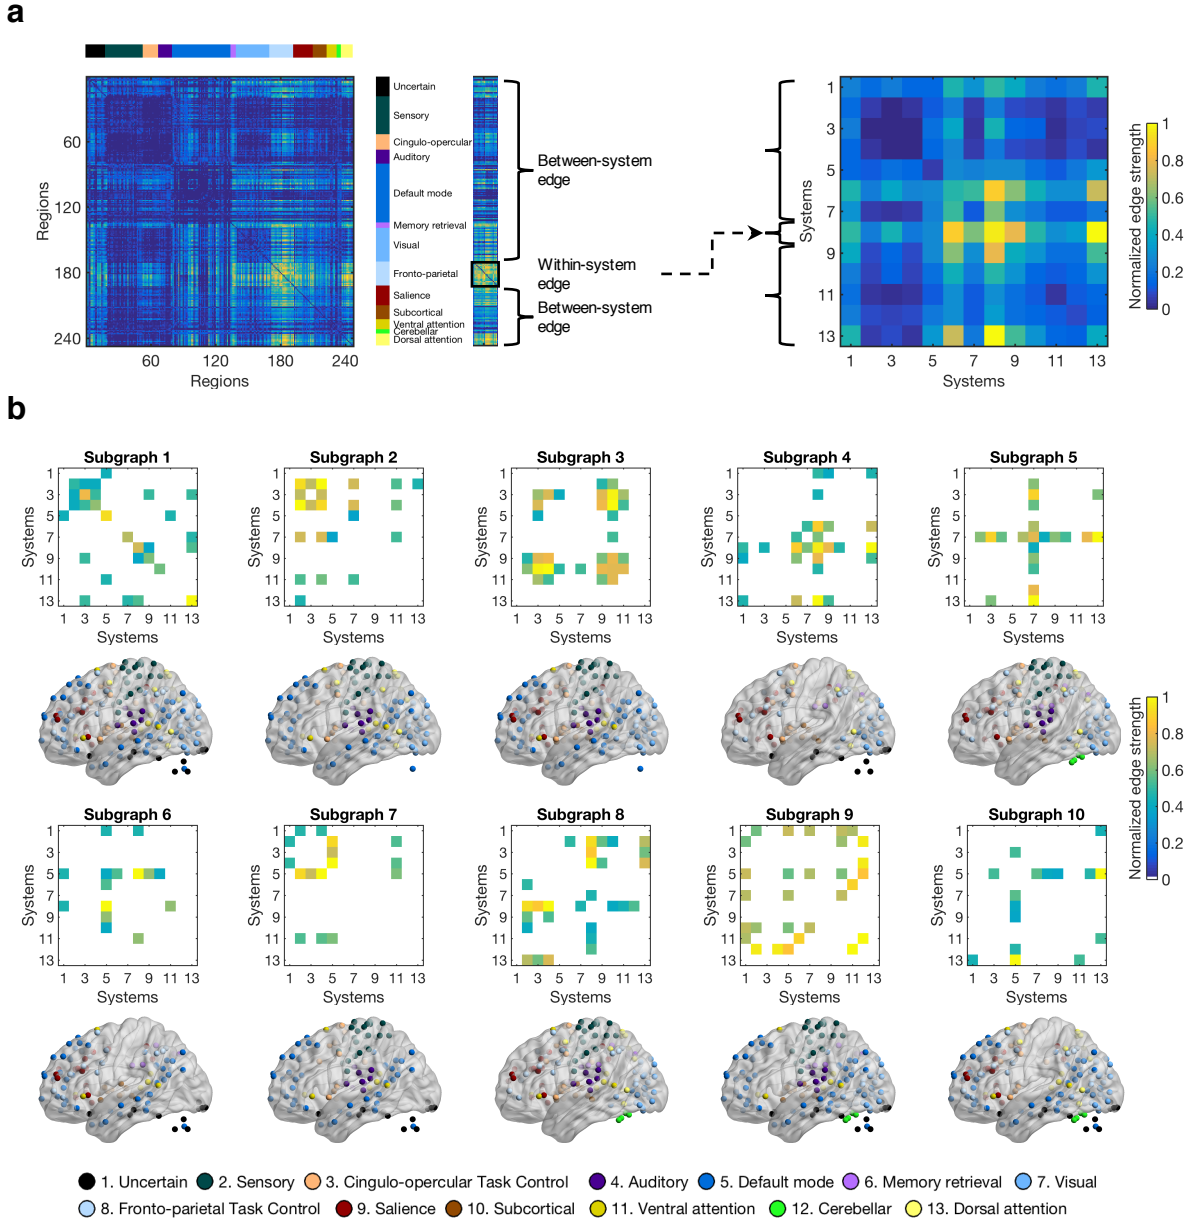


Figure 3. Patterns of connectivity in subgraphs. (a) Converting edges between nodes into edges between systems. First, the edges of each subgraph can be folded into a square matrix, representing the edges between every pair of nodes (ROIs). Then, based on the 13 putative functional systems reported by Power et al. (2011), we categorized each edge according to the system(s) to which the two nodes (ROIs) belonged. We calculated the mean strength of edges linking a node in one system to a node in another system, and refer to that value as the between-system edge. Similarly, we calculated the mean strength of edges linking two nodes that both belong to the same system and refer to that value as the within-system edge. Edges between nodes and edges between systems were normalized into the scale between 0 and 1. (b) Edges between systems in the ten subgraphs identified by NMF. For each subgraph, the top matrix shows the significant edges in that subgraph within or between systems. We show only significant edges ($p < 0.05$ after the Bonferroni correction for multiple comparisons). All nodes from systems involved in significant edges are shown on the brain below. Subgraphs varied in the degree to which they represent interactions within the same system (e.g., subgraph 1) versus interactions between different systems (e.g., subgraph 10).

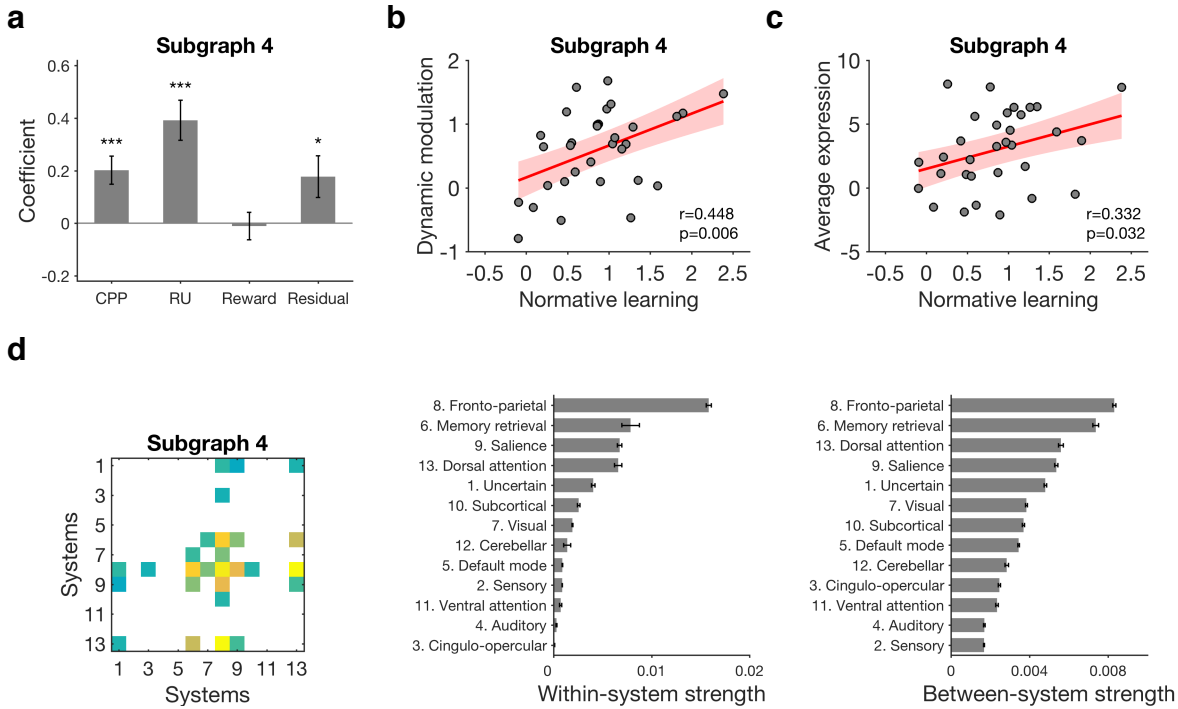


Figure 4. Temporal expression of subgraph 4 was related to learning factors and individual differences. (a) Modulation of temporal expression of subgraph 4 by learning factors. A regression model that included CPP, RU, reward and residual updating as predictors of temporal expression (indexed as the difference between positive and negative expression) of subgraph 4 was fitted for each participant, and coefficients were tested on the group level. The results showed positive effects of CPP, RU and residual updating. Error bars represent one SEM. (* $p<0.05$, *** $p<0.001$) (b) The relationship between individual normative learning and the dynamic modulation of subgraph 4 expression by normative factors. This dynamic modulation was indexed as the sum of the coefficients of CPP and RU in (a), and represents the extent to which trial-by-trial expression was influenced by the two normative learning factors. There was a significant positive correlation across participants. Each point represents one participant. The red line represents the regression line and the shaded area represents the 95% confidence interval. (c) The relationship between individual normative learning and average expression of subgraph 4. There was a significant positive correlation across participants. Each point represents one participant. The red line represents the regression line and the shaded area represents the 95% confidence interval. (d) Summary of the pattern of connectivity in subgraph 4. We summarized the pattern of connectivity as within-system strength (which is the value in the diagonal) and between-system strength (which is the average of values in the off-diagonal) for each system. The fronto-parietal system showed the strongest contributions to the sub-graphs in terms of both within-system and between-system strength. The 95% confidence interval of each system was estimated by bootstrapping 10,000 times on the edges of that system.

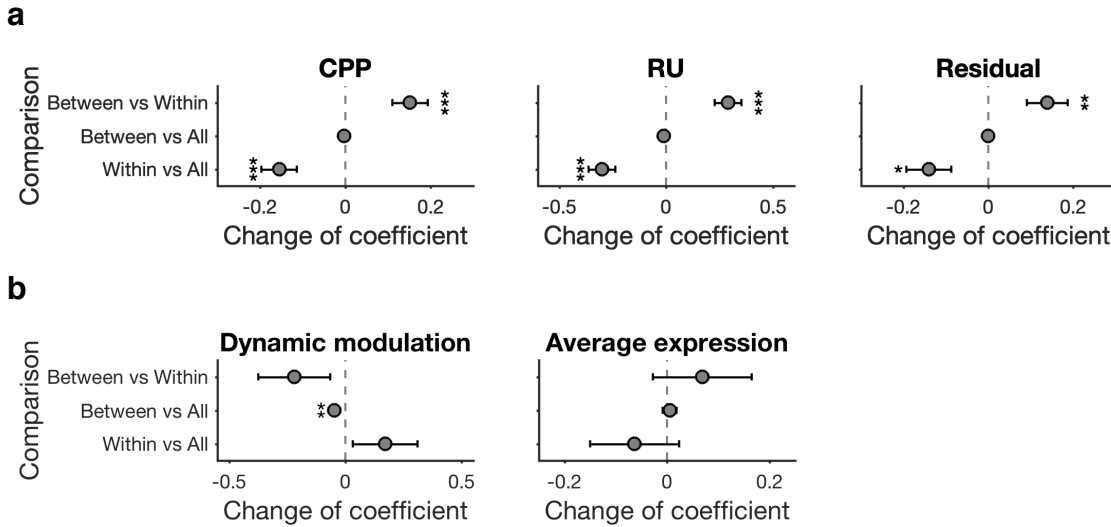


Figure 5. The contribution of between-system and within-system edges to effects of learning factors and individual differences on subgraph 4 expression. (a) The contribution of between-system and within-system edges to the effect of learning factors on temporal expression of subgraph 4. To determine the relative contribution of between- and within-system edges on time-dependent subgraph 4 expression, we performed three comparisons. First, removing between-system edges (Within versus All) decreased the effect of CPP, RU and residual updating. Second, in contrast, after removing within-system edges (Between vs All), there was no significant change in these coefficients. Third, we directly compared the effects contributed from between-system edges only and from within-system edges only (Between versus Within). For between-system edges, there were stronger positive effects for CPP, RU and residual updating. Error bars represent one SEM. (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$) (b) The contribution of between-system and within-system edges to the relationship between normative learning and dynamic modulation and average expression of subgraph 4. We performed the same three comparisons to determine the relative contribution of between- and within-system edges for each relationship with individual differences. For the effect of dynamic modulation, removing within-system edges (Between versus All) decreased the correlation coefficient. This correlation coefficient was also larger for within-system edges only than between-system edges only, but this effect was not statistically significant. For the effect of average expression, removing between-system edges (Within versus All) decreased the correlation coefficient, and the correlation coefficient was larger for between-system edges only than within-system edges only, though neither of these effects were statistically significant. Error bars represent one SEM (** $p < 0.01$).

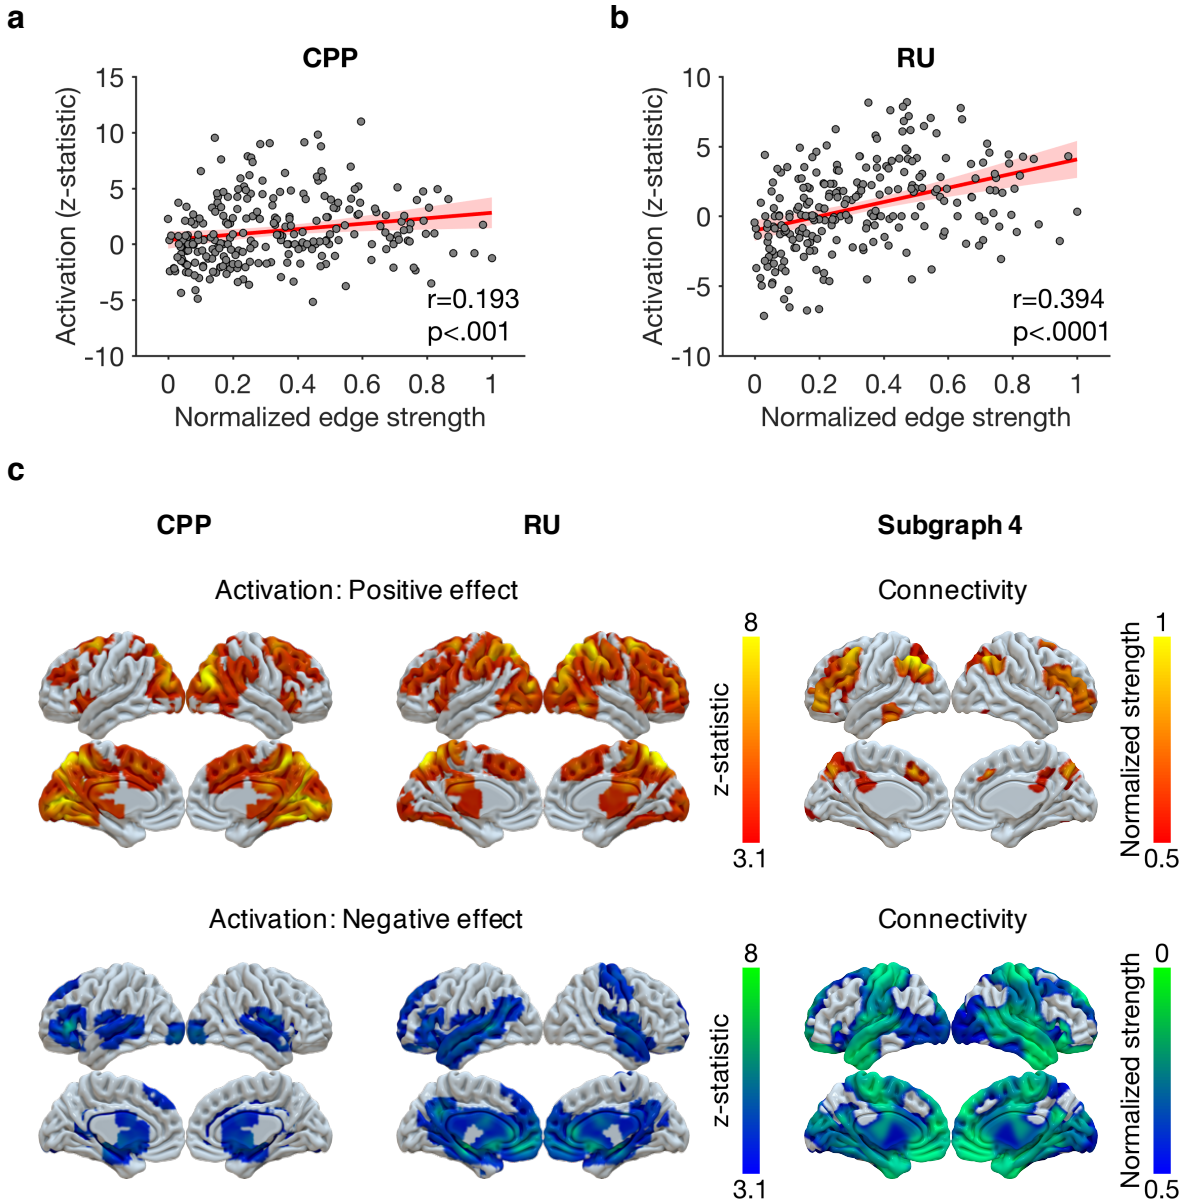
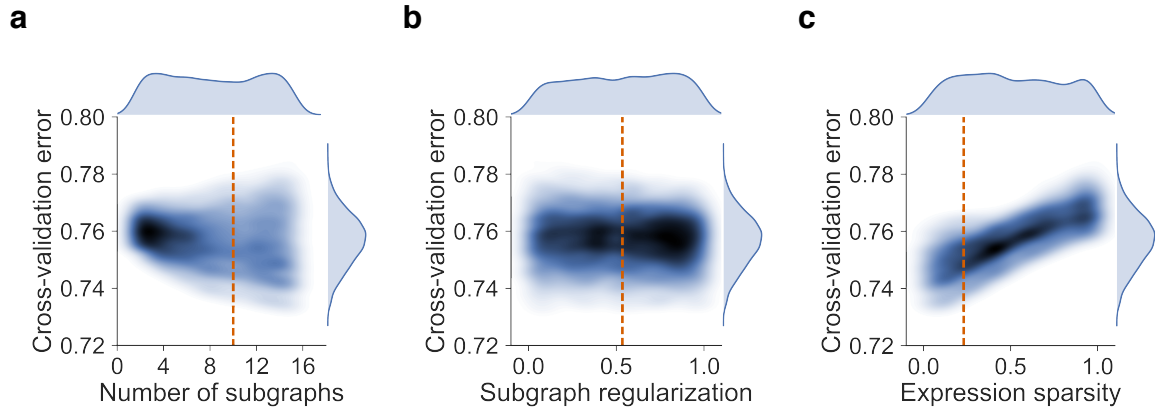
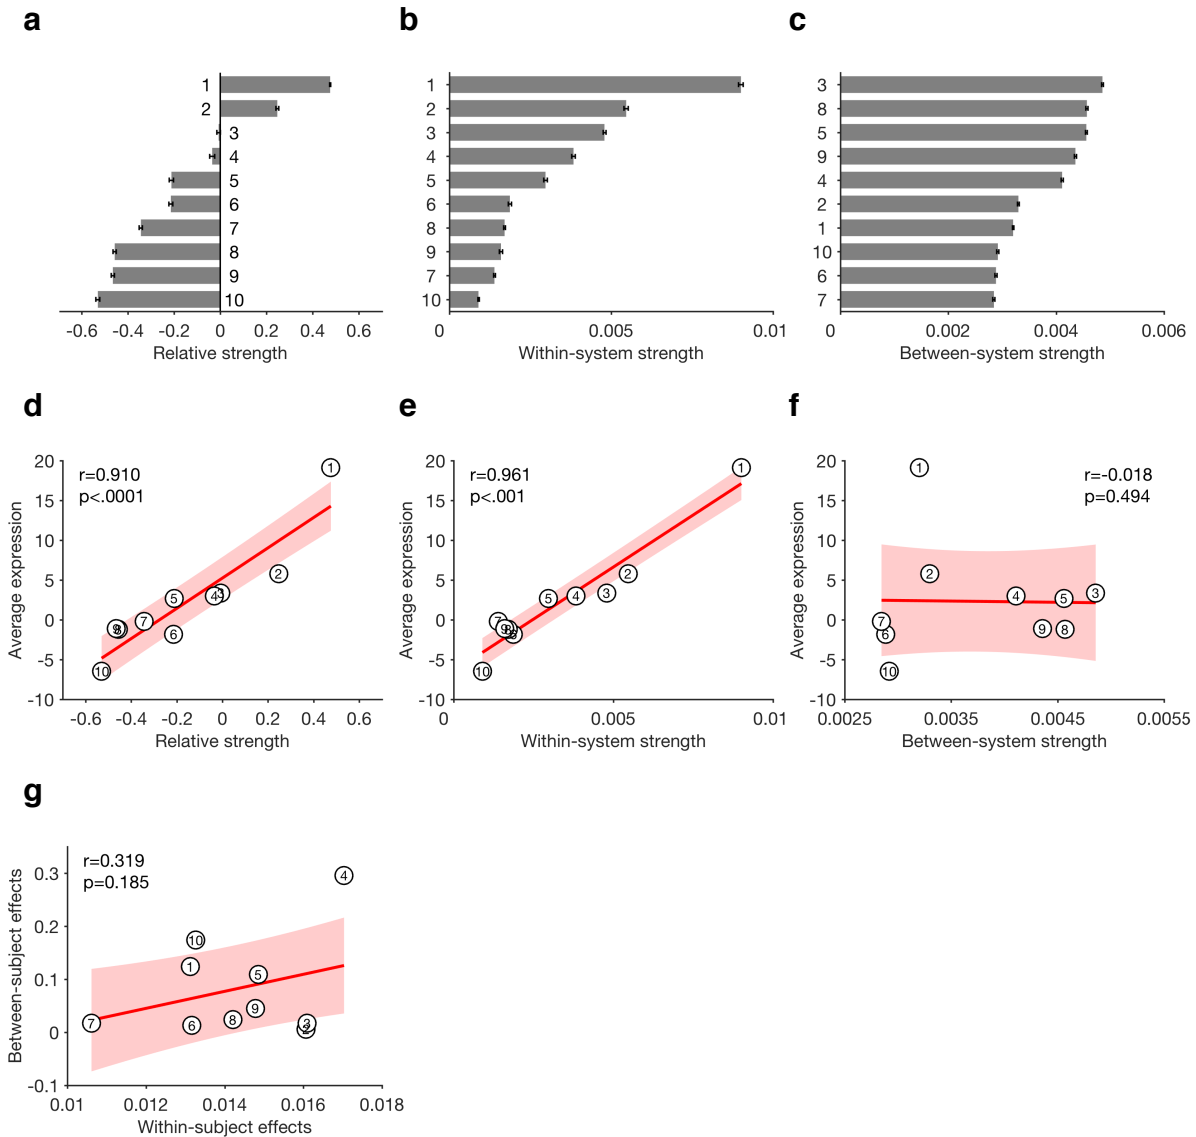


Figure 6. Relationship between edge strength of subgraph 4 and univariate task activations. (a) Relationship between the activation for CPP and the edge strength of subgraph 4. We calculated the Pearson correlation coefficient between the z -statistic for CPP from McGuire et al. (2014) and the edge strength across nodes in subgraph 4. Each data point represents an ROI. The edge strength for each ROI was calculated as the column sum of that ROI's edges to other ROIs, reflecting the summed interactions between that ROI and all others. The edges were normalized into the scale between 0 and 1. A significantly positive correlation was observed. The red line represents the regression line and the shaded area represents the 95% confidence interval. (b) Relationship between the activation for RU and the edge strength of subgraph 4. We observed a significant positive correlation between the z -statistic for RU from McGuire et al. (2014) and the edge strength across nodes in subgraph 4. The red line represents the regression line and the shaded area represents the 95% confidence interval. (c) Whole-brain thresholded activation maps for CPP and RU from McGuire et. al (2014) and whole-brain maps for edge strength of subgraph 4 in the current study.

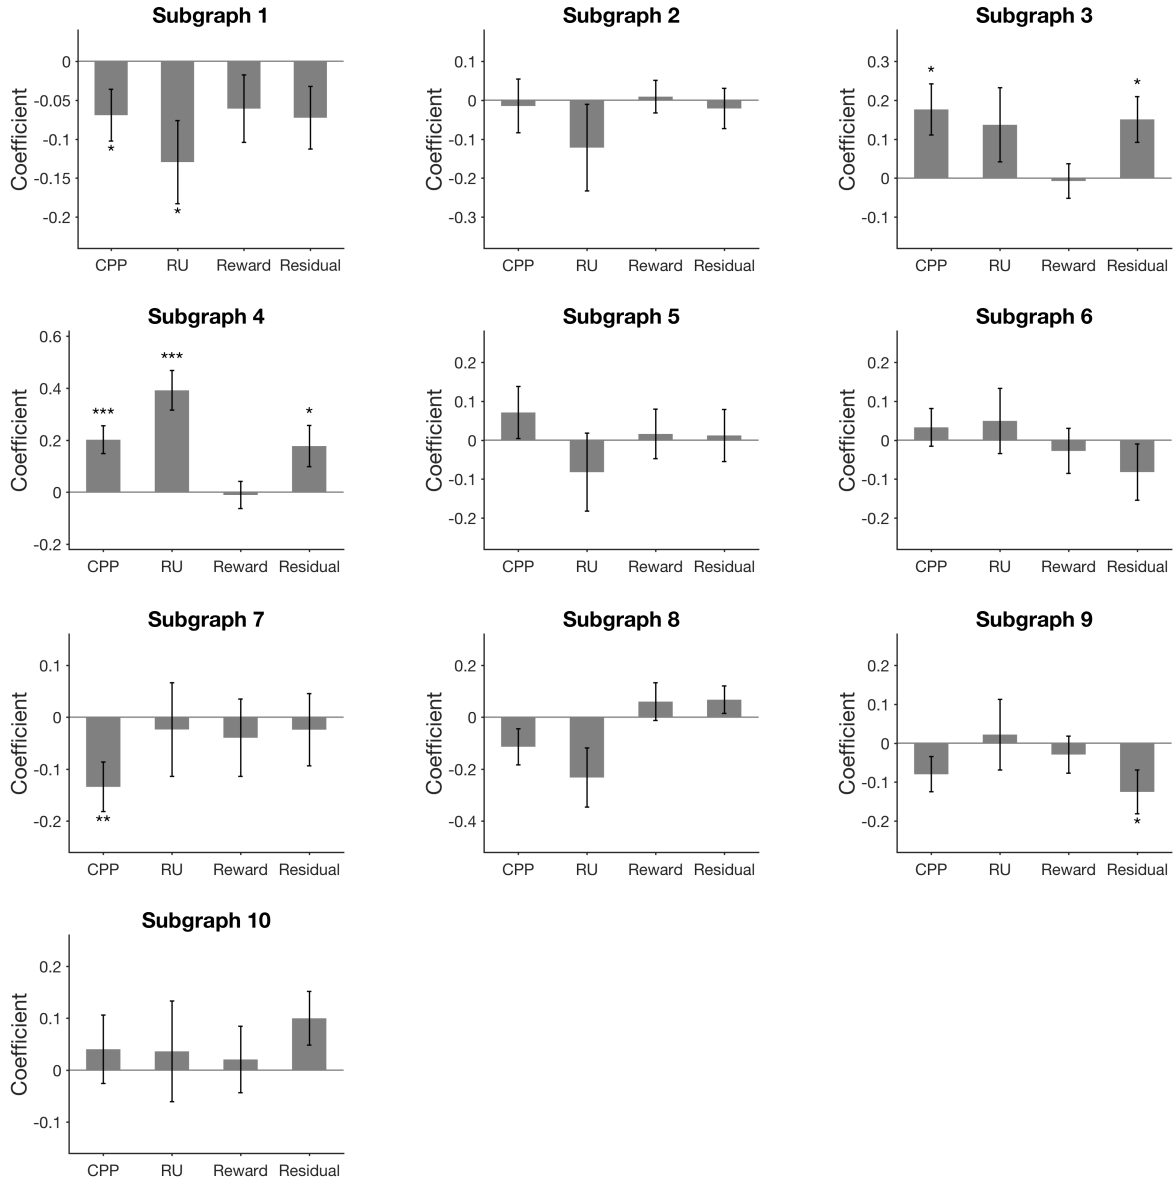


Supplementary Figure 1. Optimal parameters for nonnegative matrix factorization. (a) Number of subgraphs. We randomly sampled the number of subgraphs from a uniform distribution ($k \in [2, 15]$). The contour plot shows the Kernel density of the bivariate distribution. The darker blue area represents the higher probability mass. We selected the optimal parameter ($k = 10$) by averaging the parameter values that ensured that the cross-validation error was in the bottom 25% of the sampling distribution (orange dashed line). (b) Subgraph regularization. We randomly sampled values of subgraph regularization from a uniform distribution ($\alpha \in [0.01, 1.0]$) and select the optimal parameter ($\alpha = 0.535$). (c) Expression sparsity. We randomly sampled values of expression sparsity from a uniform distribution ($\beta \in [0.01, 1.0]$) and select the optimal parameter ($\beta = 0.230$).

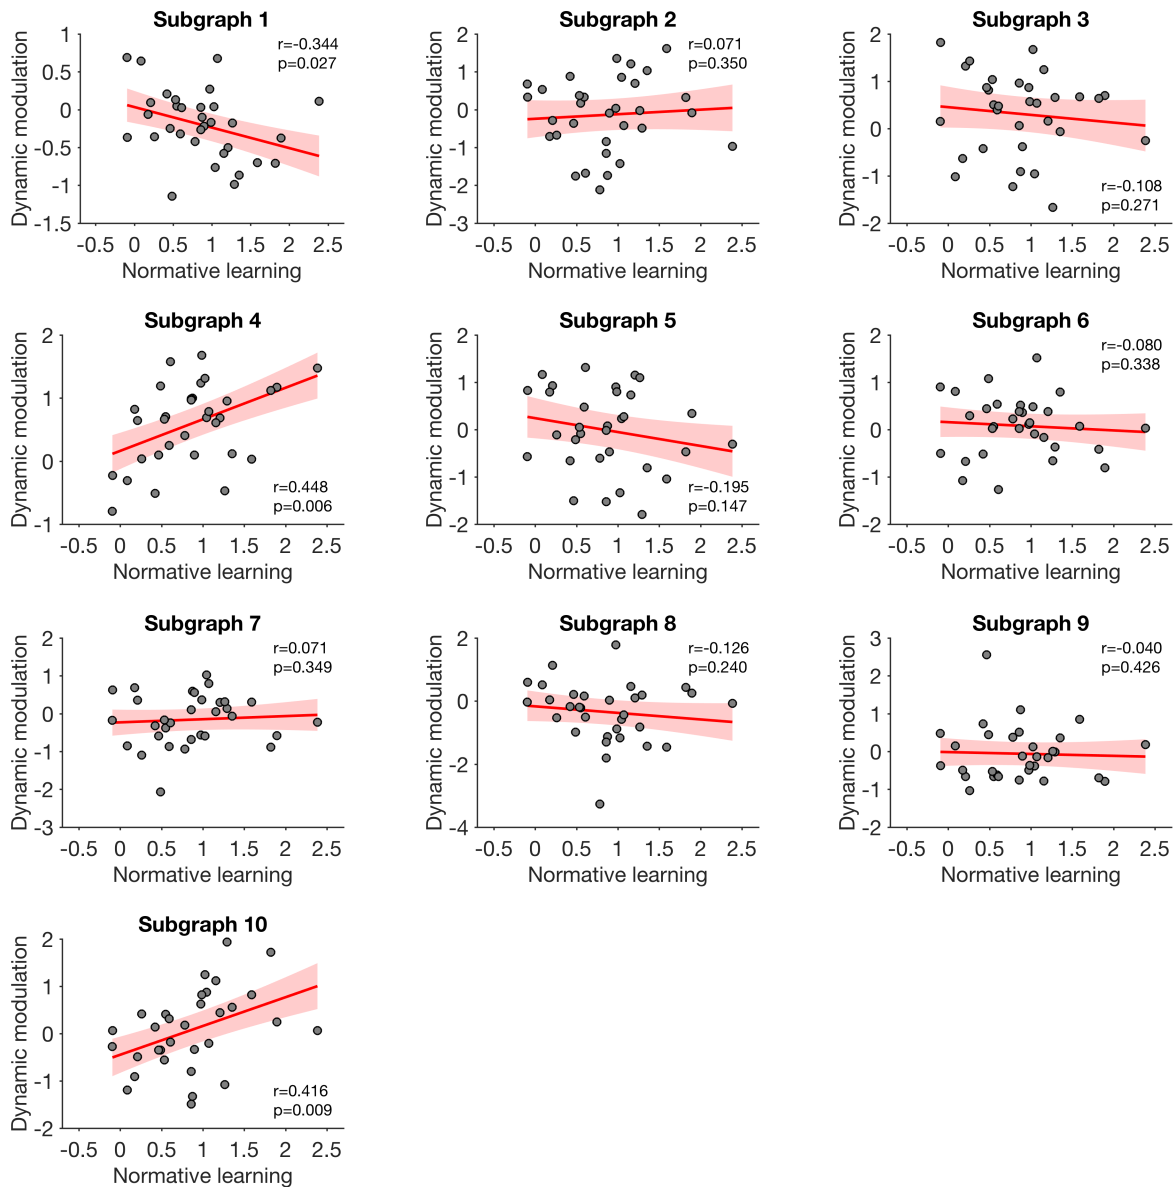


Supplementary Figure 2. Properties of subgraphs. (a) Subgraphs differ in the extent of within- versus between-system edge strength. For each subgraph, the strength of within-system edges (edges linking two nodes that both belong to the same system; Fig. 3b) was averaged, and the strength of between-system edges (edges linking one node from one system to another node from another system; Fig. 3b) was averaged. The ten subgraphs are ordered according to the relative strength of within- versus between-system edges. To form a normalized relative strength, we subtracted the average strength of between-system edges from the average strength of within-system edges and then divided this difference by their sum. A high relative strength means that a subgraph has stronger within-system edges than between-system edges (e.g., subgraph 1). The 95% confidence interval of each subgraph was estimated by bootstrapping 10,000 times on the edges of that subgraph. (b) Subgraphs differ in the extent of within-system strength. For each subgraph, the strength of within-system edges was averaged. For demonstration, the ten subgraphs are ordered according to within-system strength. The 95% confidence interval of each subgraph was estimated by bootstrapping 10,000 times on the edges of that subgraph. (c) Subgraphs differ in the extent of between-system strength. For each subgraph, the strength of between-system edges was averaged. For demonstration, the ten subgraphs are ordered according to the between-system strength. The 95% confidence interval of each subgraph was estimated by bootstrapping 10,000 times on the edges of that subgraph. (d) The relationship between relative strength and average expression across subgraphs. Average expression was calculated as the difference between positive

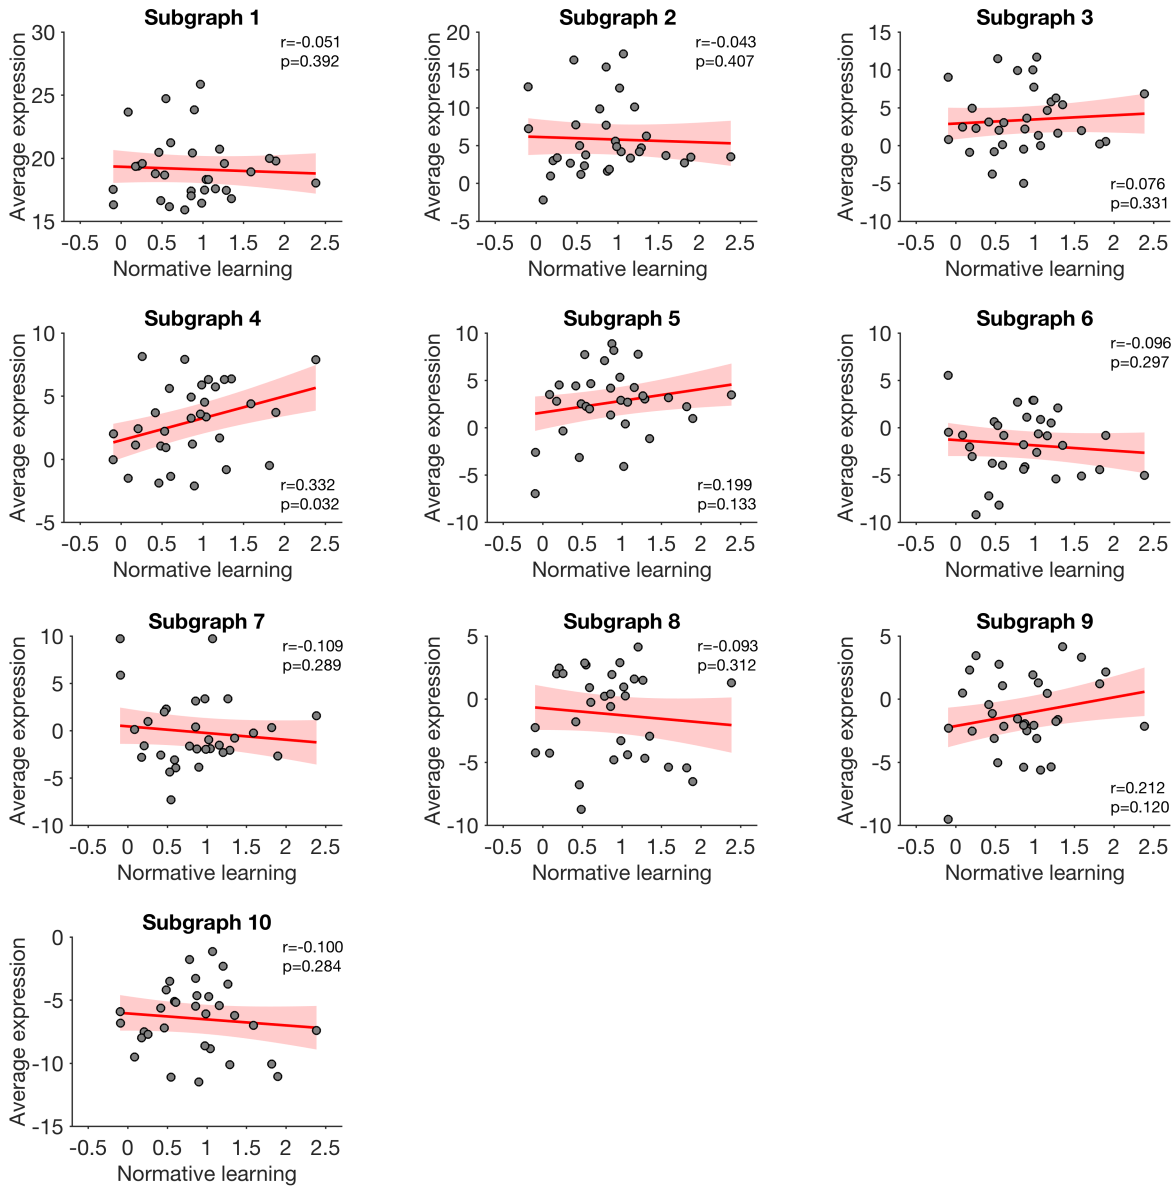
expression and negative expression. Each data point represents one subgraph. A significantly positive correlation was observed. The red line represents the regression line and the shaded area represents the 95% confidence interval. (e) The relationship between within-system strength and average expression across subgraphs. A significant positive correlation was observed. The red line represents the regression line and the shaded area represents the 95% confidence interval. (f) The relationship between between-system strength and average expression across subgraphs. There was no significant correlation. The red line represents the regression line and the shaded area represents the 95% confidence interval. (g) The relationship between explained variance of within-subject effects and explained variance of between-subject effects across subgraphs. For the within-subject effects, we implemented a regression model including factors of CPP, RU, reward and residual updating for each subgraph. Explained variance was indexed as R^2 of the regression model. For the between-subject effects, we implemented a regression model to predict individual normative learning for each subgraph. We included two regressors: dynamic modulation of normative factors (CPP and RU) on subgraph expression, and average subgraph expression. We then calculated R^2 of the regression model for each subgraph. Among the ten subgraphs, subgraph 4 showed the strongest R^2 for both within-subject and between-subject effects.



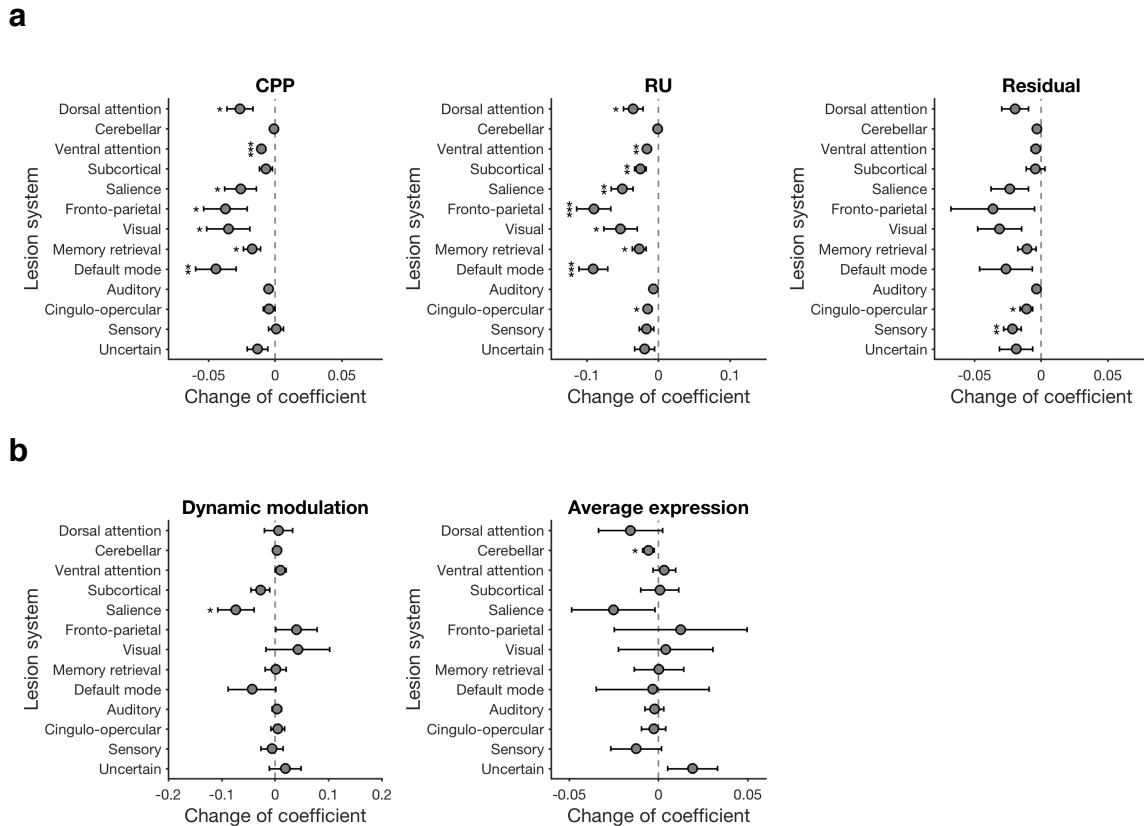
Supplementary Figure 3. Modulation of temporal expression by learning factors. A regression model that included CPP, RU, reward and residual updating as predictors of trial-by-trial expression was fitted for each participant and subgraph, and then regression coefficients were tested on the group level. Error bars represent one SEM. (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$)



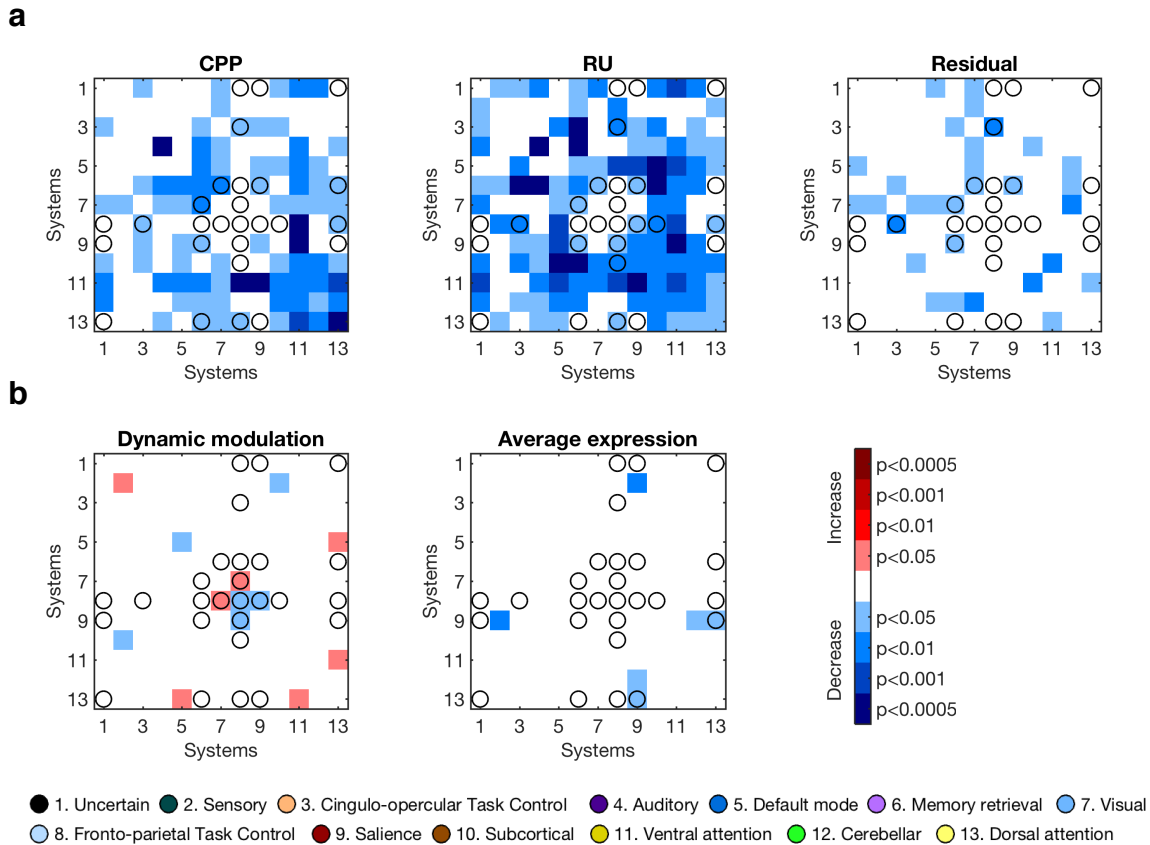
Supplementary Figure 4. The relationship between individual normative learning and the dynamic modulation of subgraph expression by normative factors. Normative learning was indexed by the sum of the coefficients of CPP and RU from a behavioral regression model, and represents the extent to which a participant’s behavior was influenced by the two normative learning factors. Dynamic modulation was indexed by the sum of coefficients of CPP and RU from the regression model against trial-by-trial expression in Supplementary Fig. 3, and represents the extent to which subgraph expression in that participant was influenced by the two normative learning factors. Each point represents one participant. The red line represents the regression line and the shaded area represents the 95% confidence interval.



Supplementary Figure 5. The relationship between individual normative learning and the average expression of each subgraph. Each point represents one participant. The red line represents the regression line and the shaded area represents the 95% confidence interval.

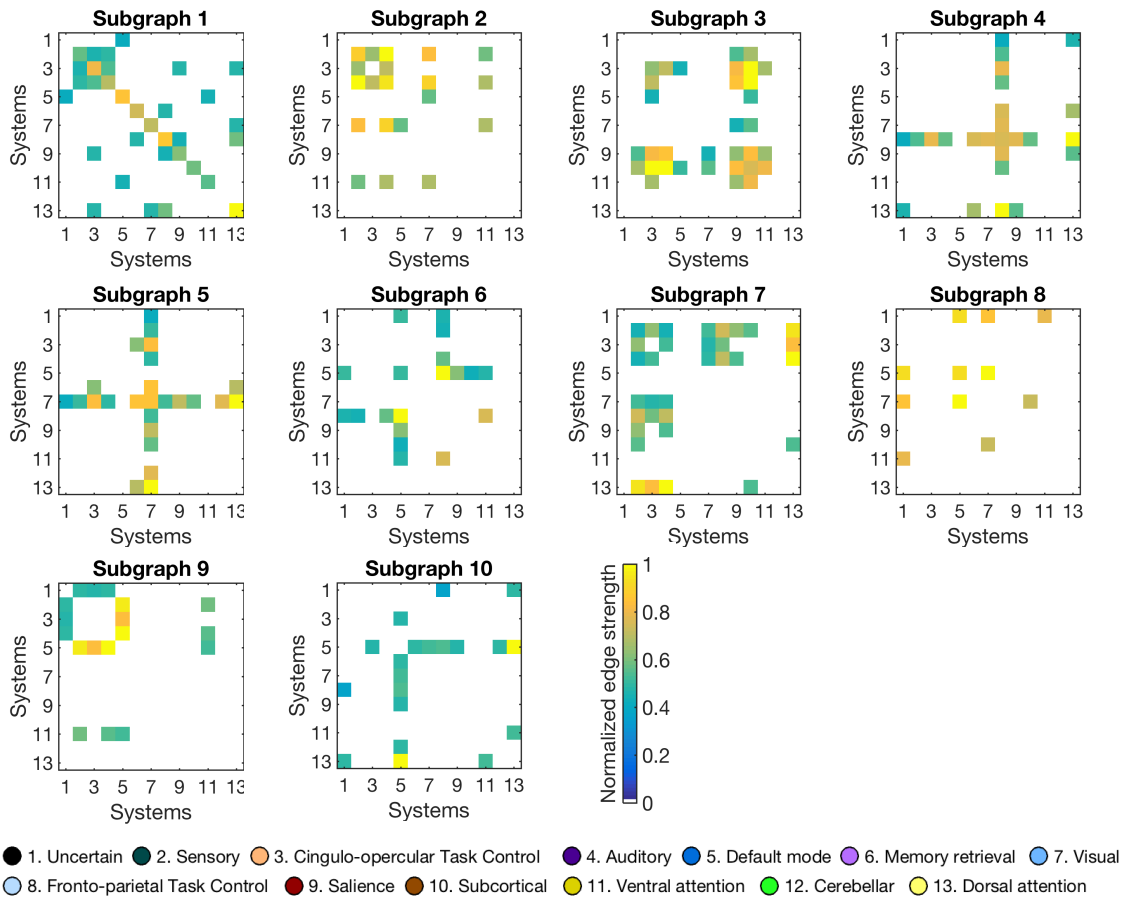


Supplementary Figure 6. Contributions of different functional systems in subgraph 4. (a) Contributions of different functional systems to the effect of learning factors on temporal expression of subgraph 4. We removed all edges from one of the 13 systems and re-estimated the coefficients for CPP, RU, reward and residual updating. Then we compared these coefficients with the original coefficients (including all the edges) to estimate the contribution of each system. Error bars represent one SEM. ($*p < 0.05$, $**p < 0.01$, $***p < 0.001$) (b) Contributions of different functional systems to the relationship between normative learning and dynamic modulation and average expression of subgraph 4. We repeated the same procedure and estimated the change of correlation coefficients for each relationship separately. Error bars represent one SEM. ($*p < 0.05$).




Supplementary Figure 7. Contributions of different system edges in subgraph 4. (a) Contributions of different system edges to the effect of learning factors on temporal expression of subgraph 4. We removed all edges for one of the 91 system-by-system connections and re-estimated the coefficients for CPP, RU, reward and residual updating. Then we compared these coefficients with the original coefficients (including all the edges) to estimate the contribution of each system edge. The open circles denote the significant system edges in subgraph 4 (as shown in Fig. 3b). An increase in coefficients is shown in red while a decrease is shown in blue. Lower p values are shown in darker color. A p value around 0.0005 corresponds to a corrected p value of .05 after multiple comparisons (i.e., $0.05/91$). (b) Contributions of different system edges to the relationship between normative learning and dynamic modulation and average expression of subgraph 4. We repeated the same procedure and estimated the change of correlation coefficients for each relationship separately.

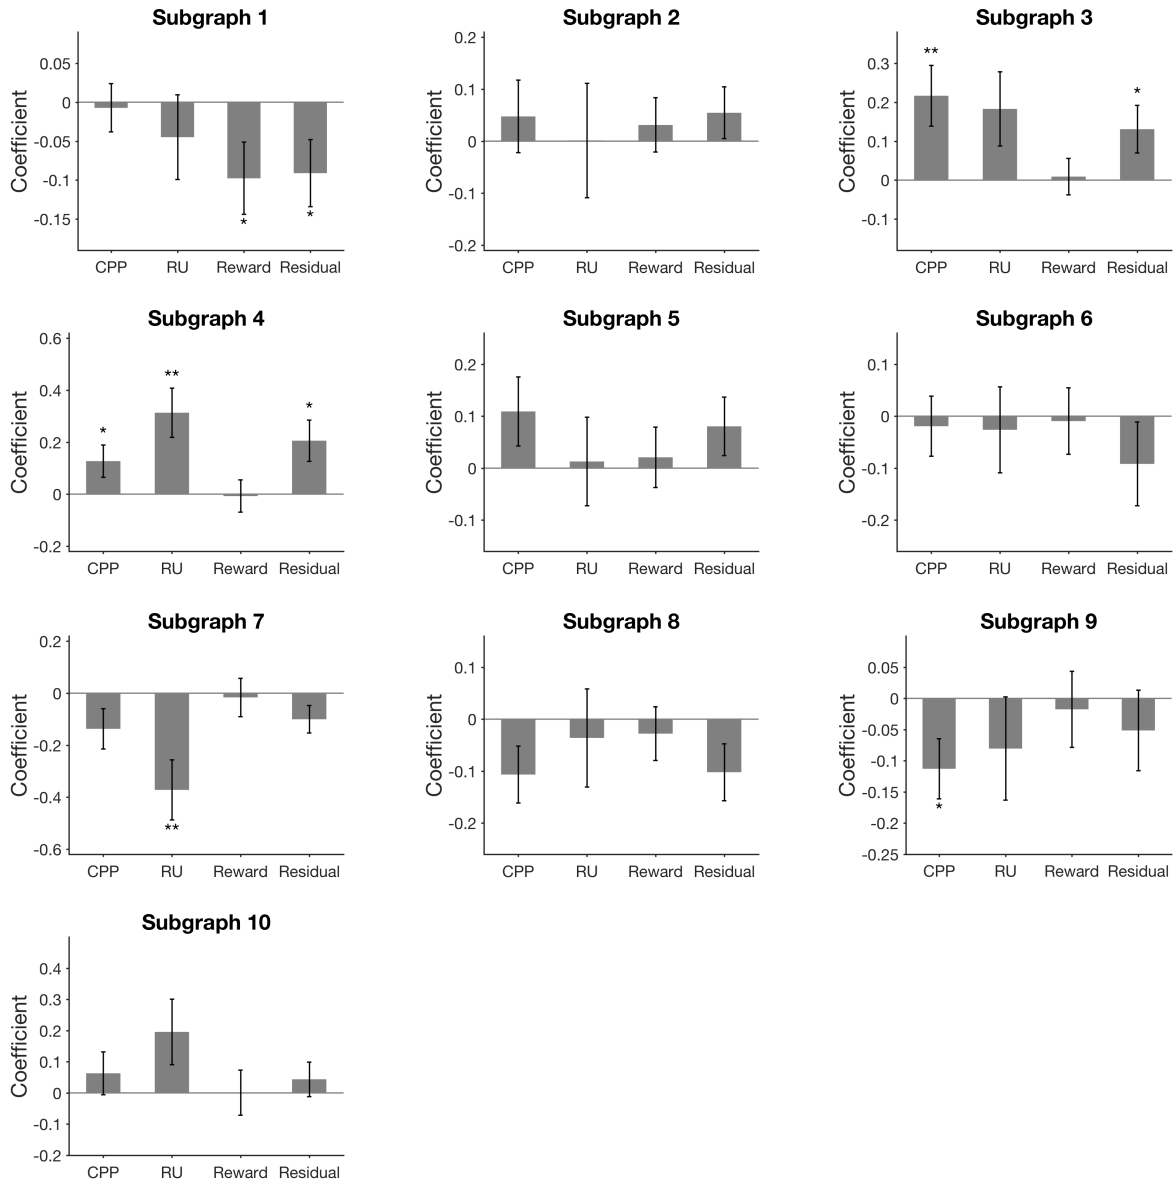
a



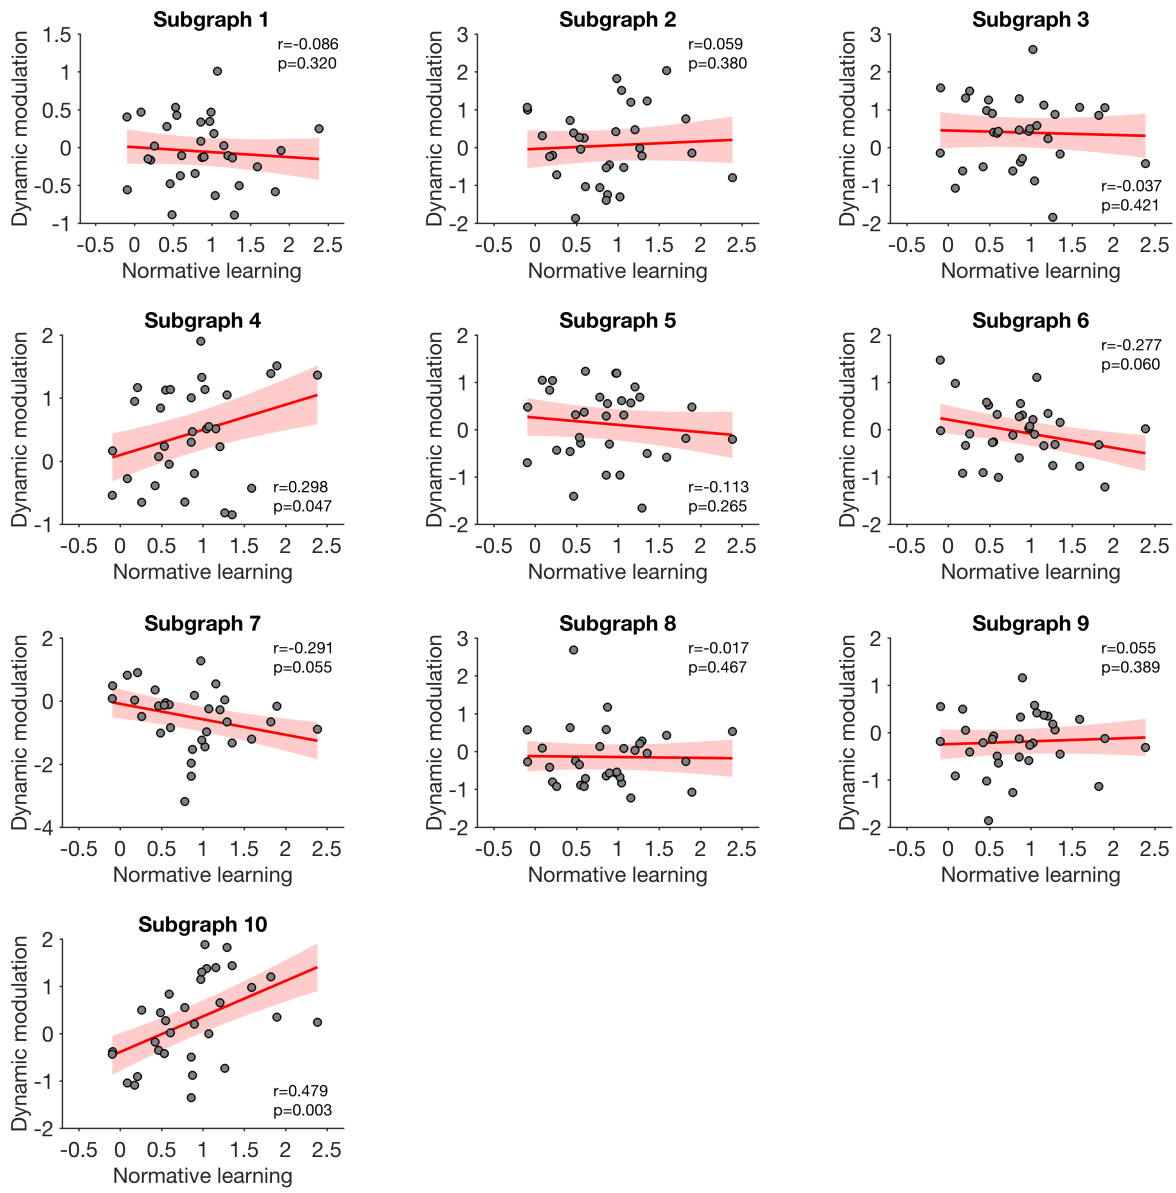
b



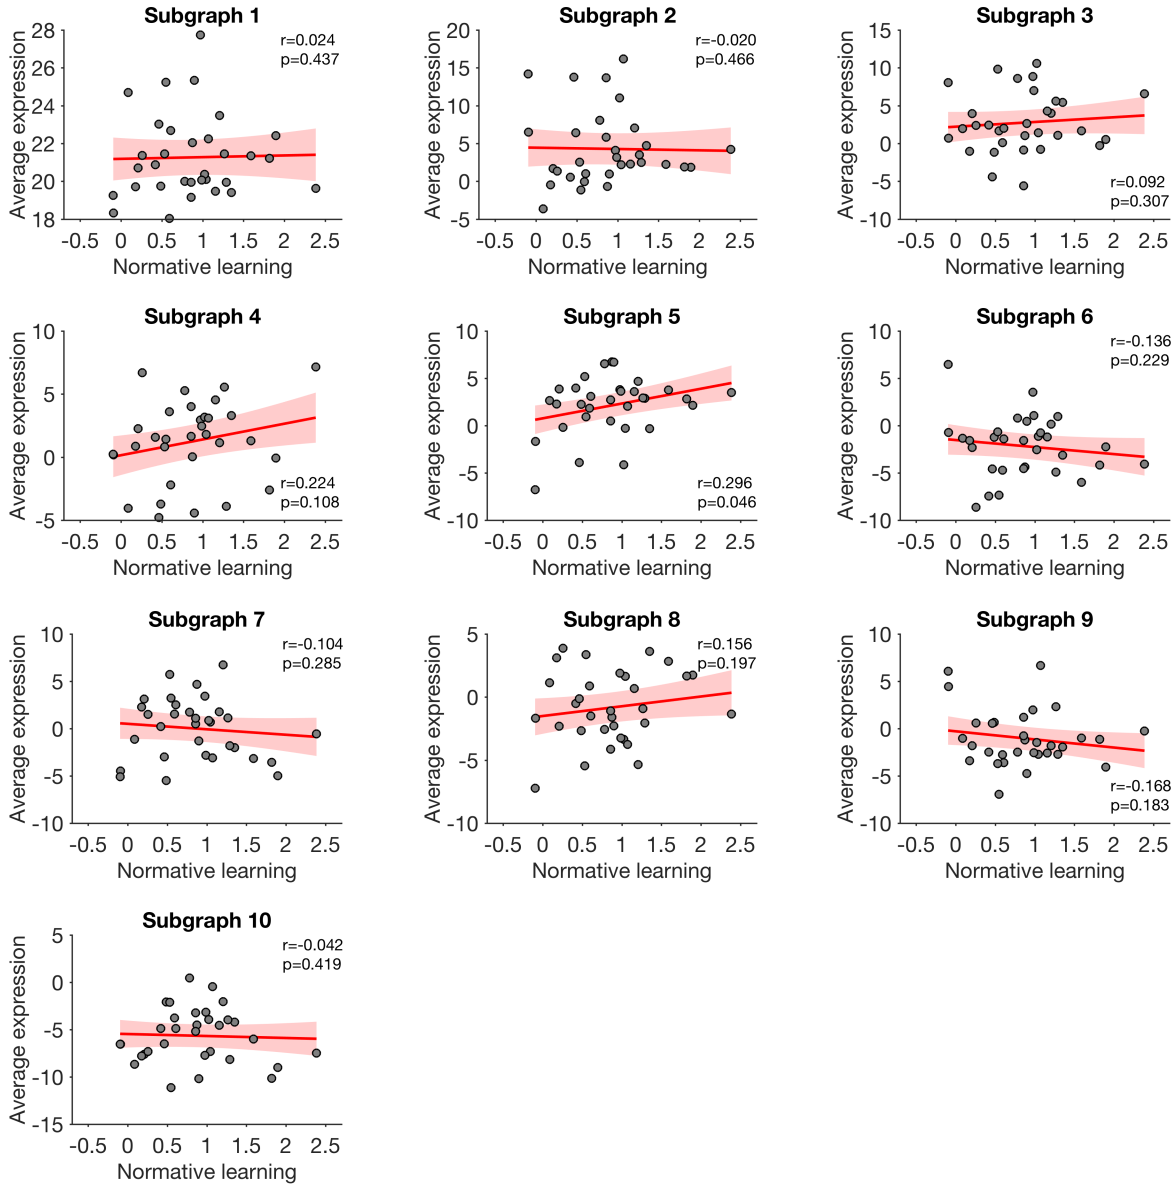
Supplementary Figure 8. Robustness check. Subgraphs identified with smaller sliding time window of 8 TRs, with 6 TRs overlapping between time windows. (a) Edges between systems in the ten subgraphs identified by NMF, as in Fig. 3b. (b) Summary of the pattern of connectivity in subgraph 4, as in Fig. 4d, showing the within-system strength and between-system strength of each functional system. The 95% confidence interval of each system was estimated by bootstrapping 10,000 times on the edges of that system.



Supplementary Figure 9. Robustness check. Modulation of temporal expression by learning factors in all ten subgraphs identified with smaller sliding time window of 8 TRs (Supplementary Fig. 8). Error bars represent one SEM. (* $p < 0.05$, ** $p < 0.01$)

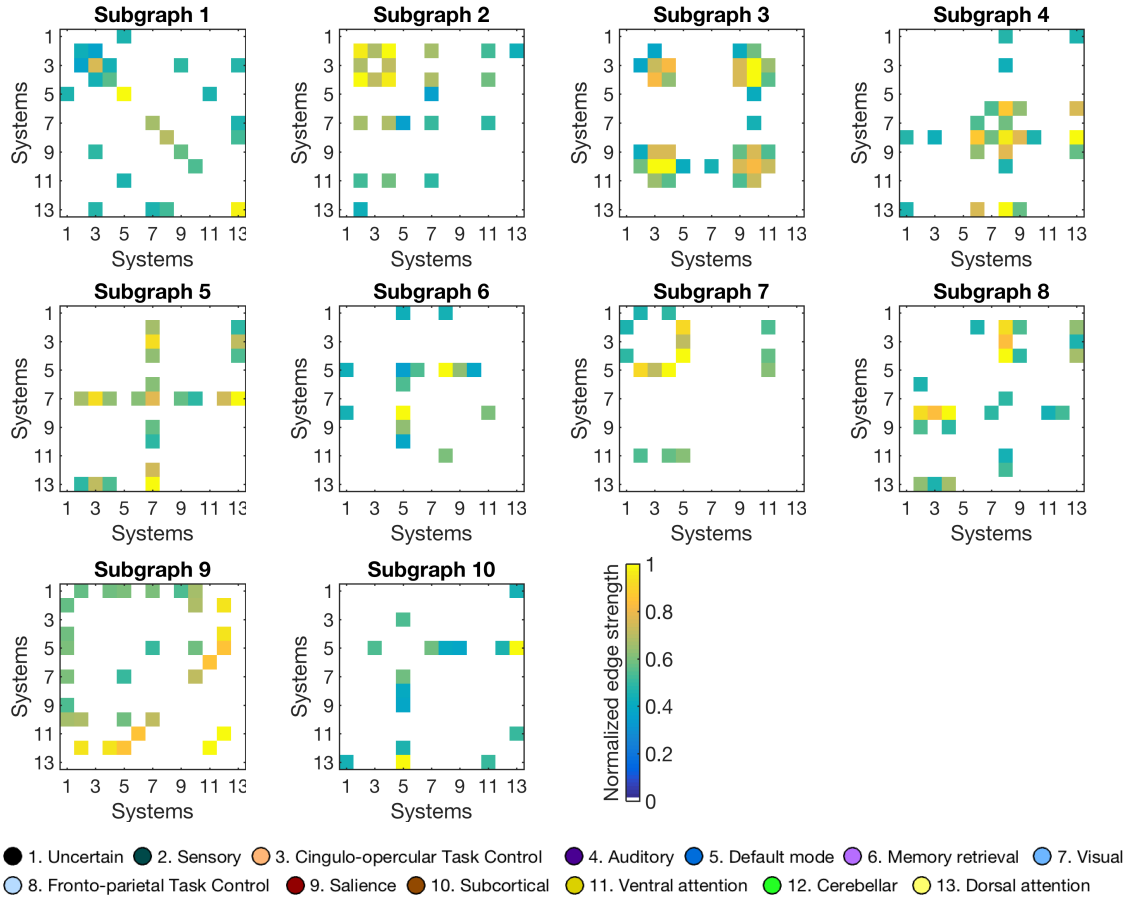


Supplementary Figure 10. Robustness check. The relationship between individual normative learning and the dynamic modulation of subgraph expression by normative factors in all ten subgraphs identified with smaller sliding time window of 8 TRs (Supplementary Fig. 8). Each point represents one participant. The red line represents the regression line and the shaded area represents the 95% confidence interval.




Supplementary Figure 11. Robustness check. The relationship between individual normative learning and the average subgraph expression in all ten subgraphs identified with smaller sliding time window of 8 TRs (Supplementary Fig. 8). Each point represents one participant. The red line represents the regression line and the shaded area represents the 95% confidence interval.

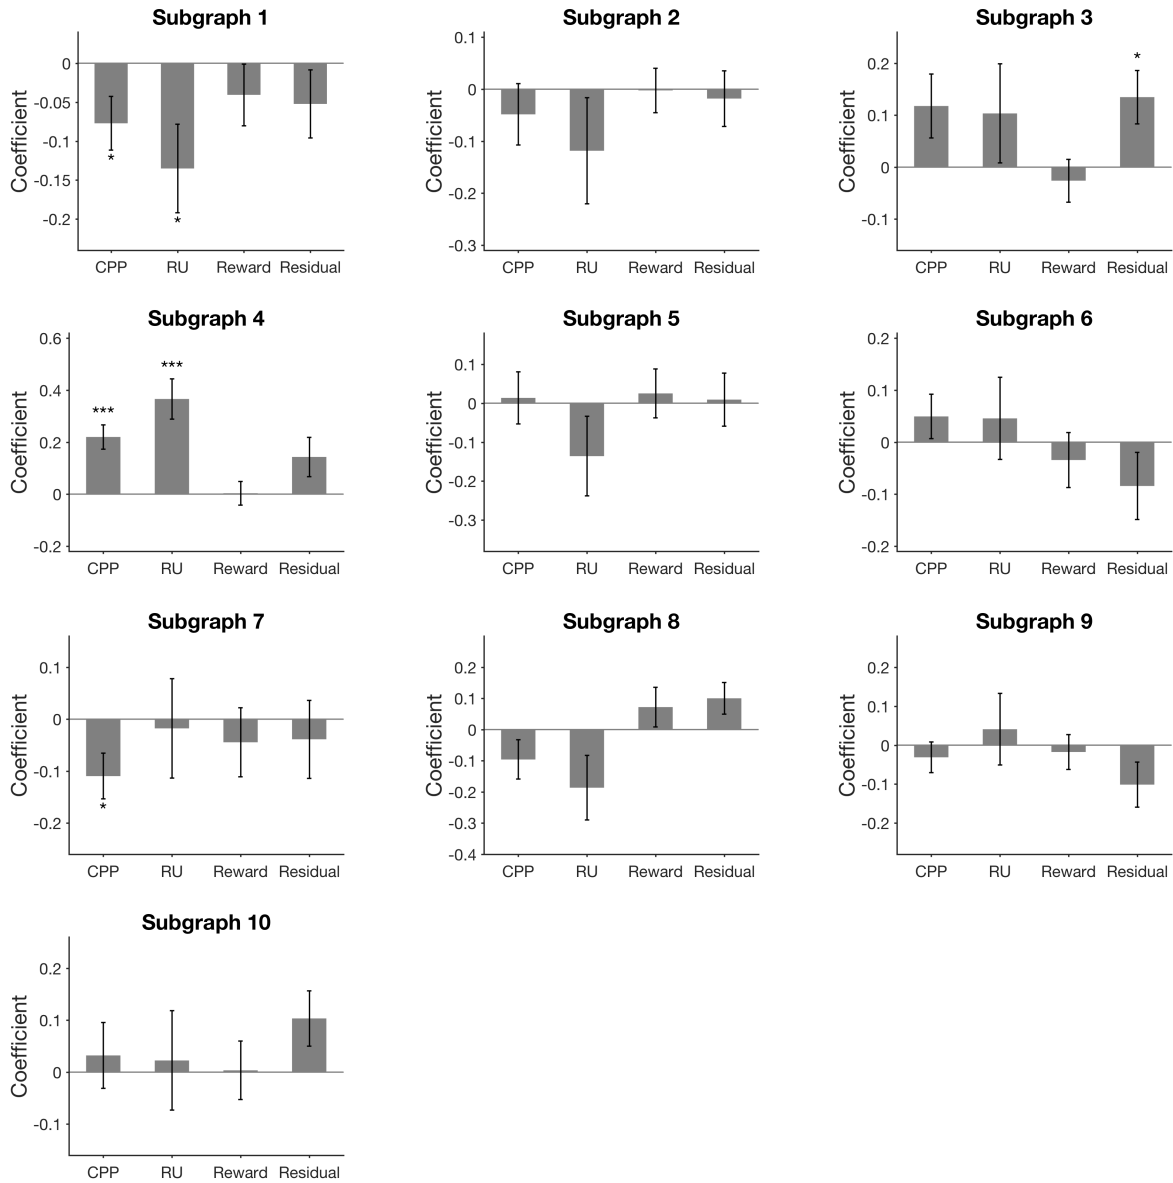
a



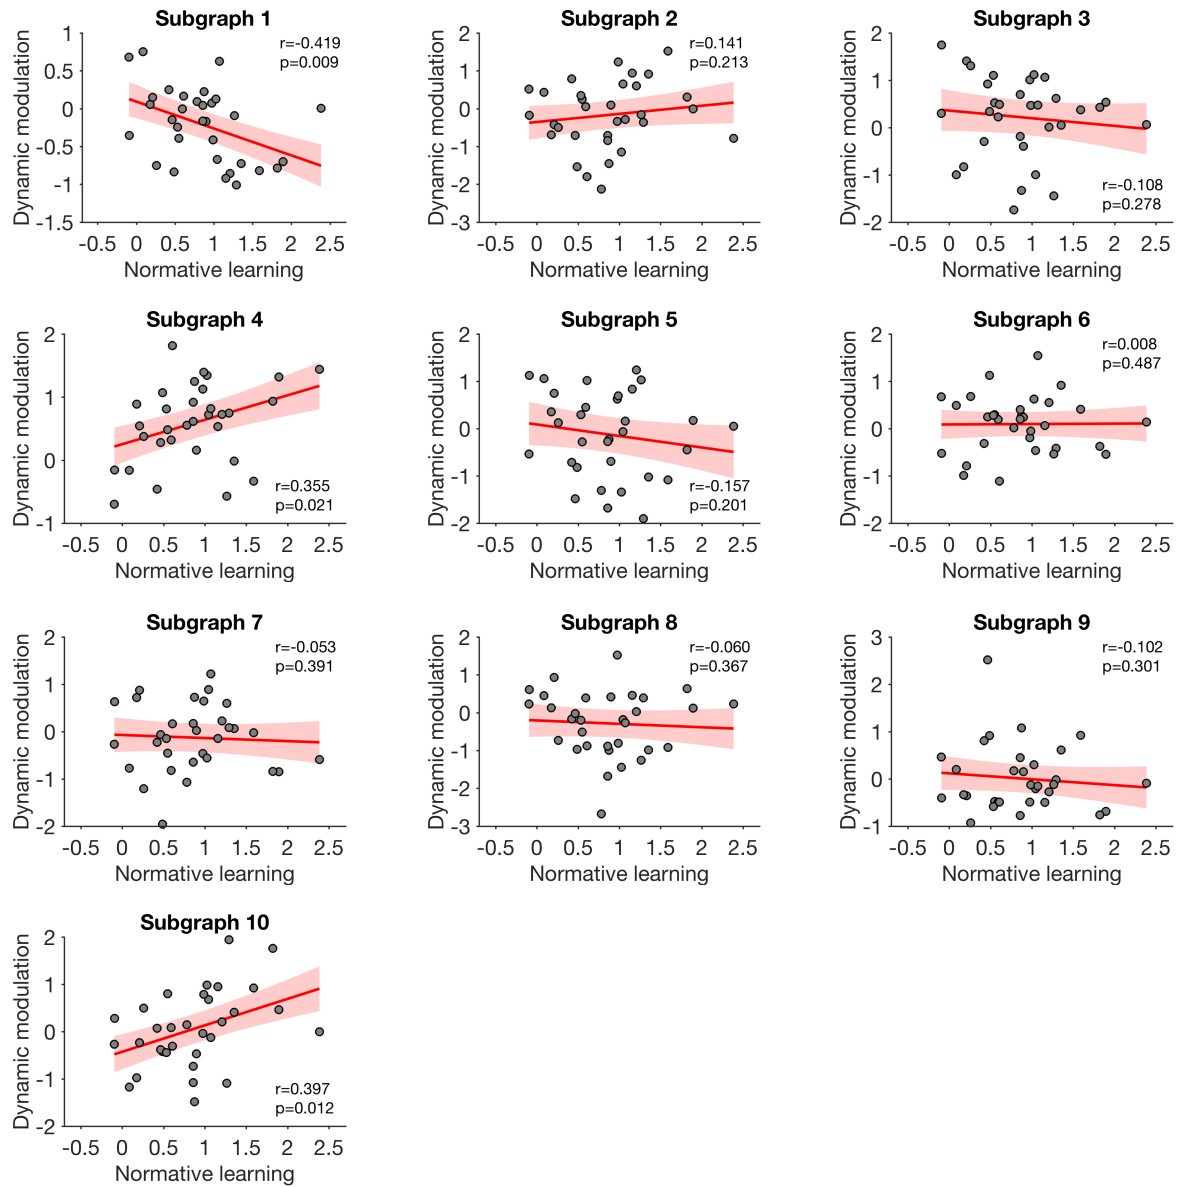
b



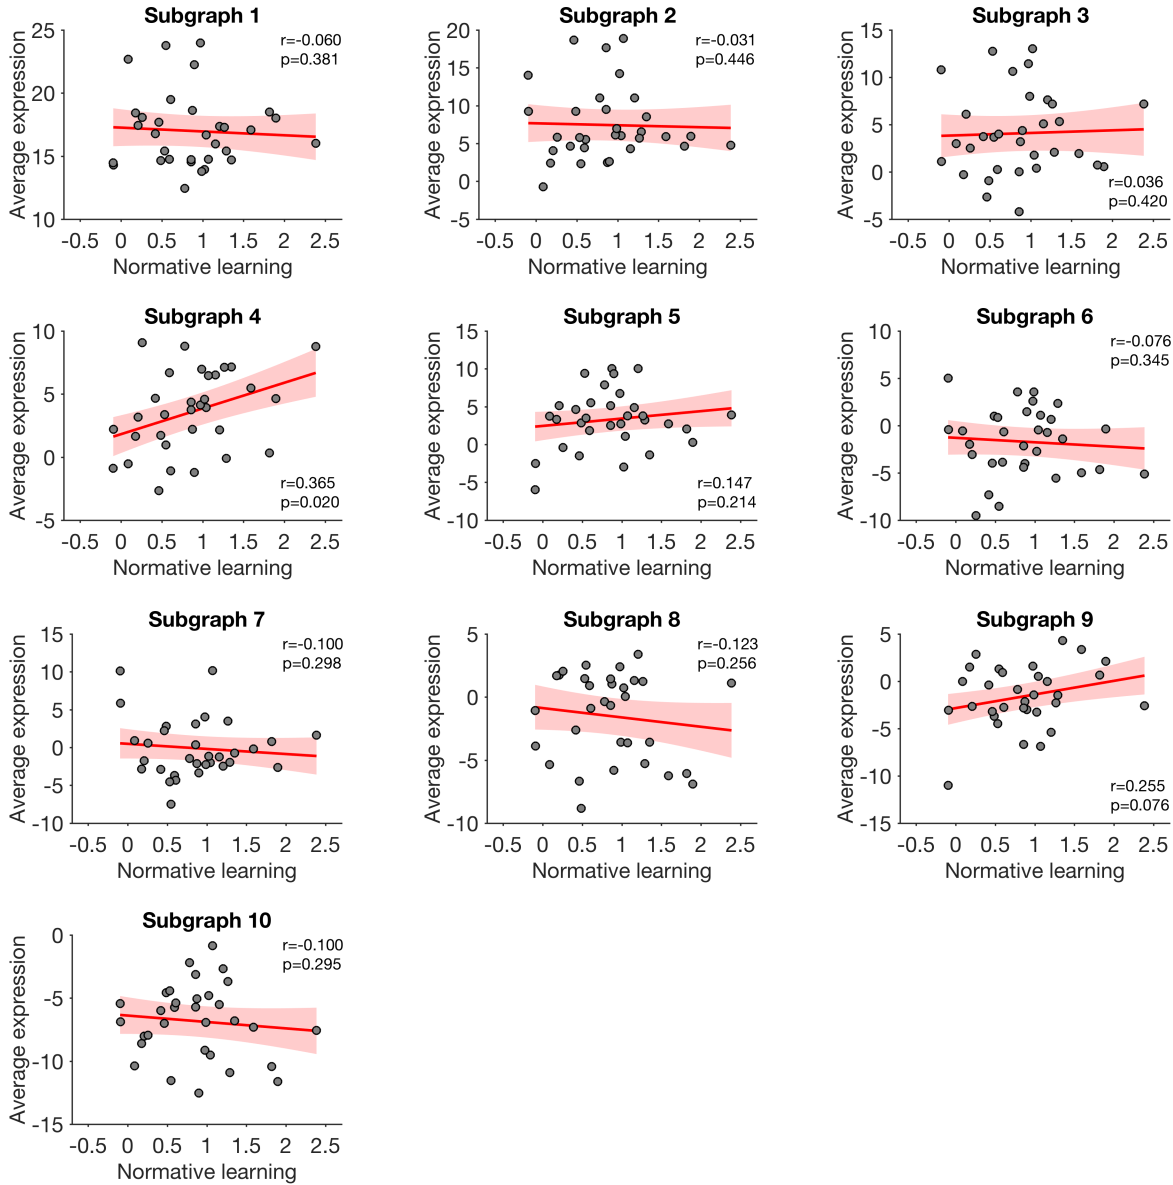
Supplementary Figure 12. Robustness check. Subgraphs identified with larger sliding time window of 12 TRs, with 10 TRs overlapping between time windows. (a) Edges between systems in the ten subgraphs identified by NMF, as in Fig. 3b. (b) Summary of the pattern of connectivity in subgraph 4, as in Fig. 4d, showing the within-system strength and between-system strength of each functional system. The 95% confidence interval of each system was estimated by bootstrapping 10,000 times on the edges of that system.



Supplementary Figure 13. Robustness check. Modulation of temporal expression by learning factors in all ten subgraphs identified with larger sliding time window of 12 TRs (Supplementary Fig. 12). Error bars represent one SEM. (* $p < 0.05$, *** $p < .001$)

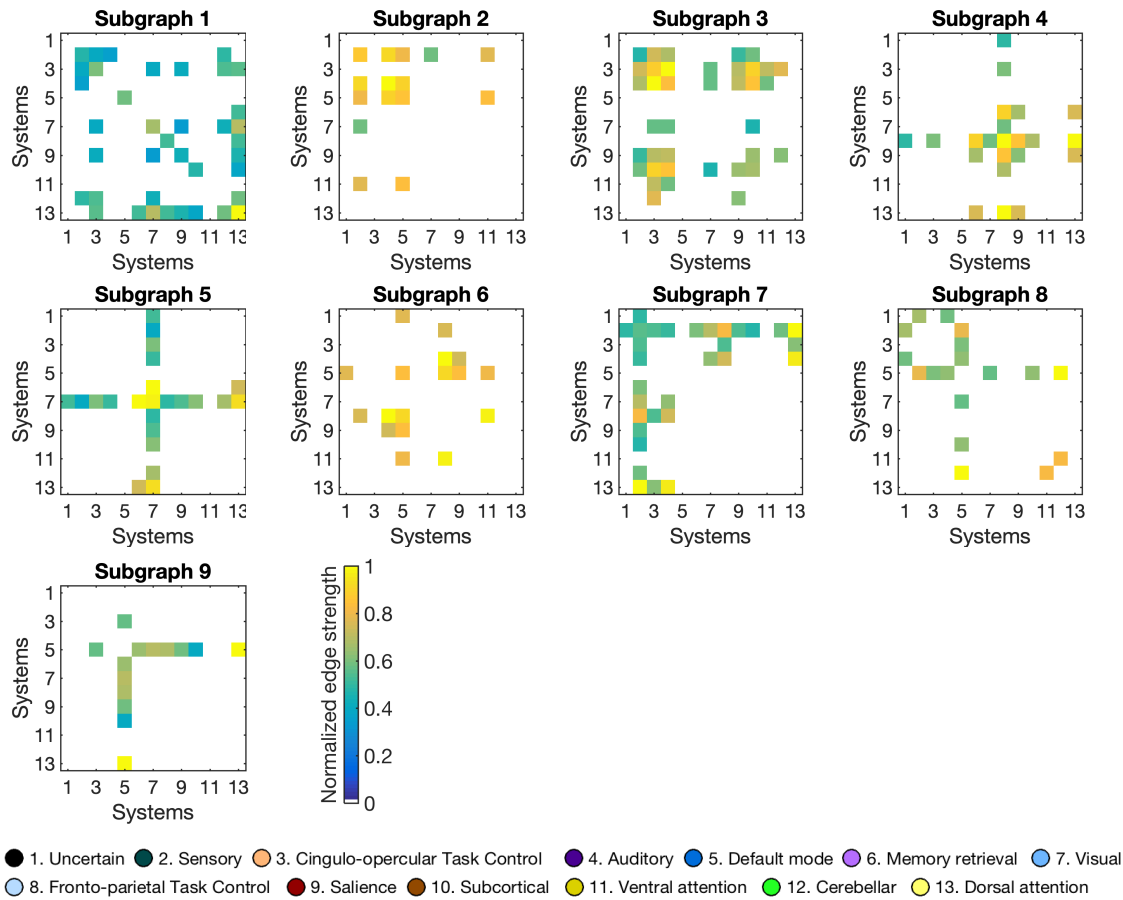


Supplementary Figure 14. Robustness check. The relationship between individual normative learning and the dynamic modulation of subgraph expression by normative factors in all ten subgraphs identified with larger sliding time window of 12 TRs (Supplementary Fig. 12). Each point represents one participant. The red line represents the regression line and the shaded area represents the 95% confidence interval.




Supplementary Figure 15. Robustness check. The relationship between individual normative learning and the average subgraph expression in all ten subgraphs identified with larger sliding time window of 12 TRs (Supplementary Fig. 12). Each point represents one participant. The red line represents the regression line and the shaded area represents the 95% confidence interval.

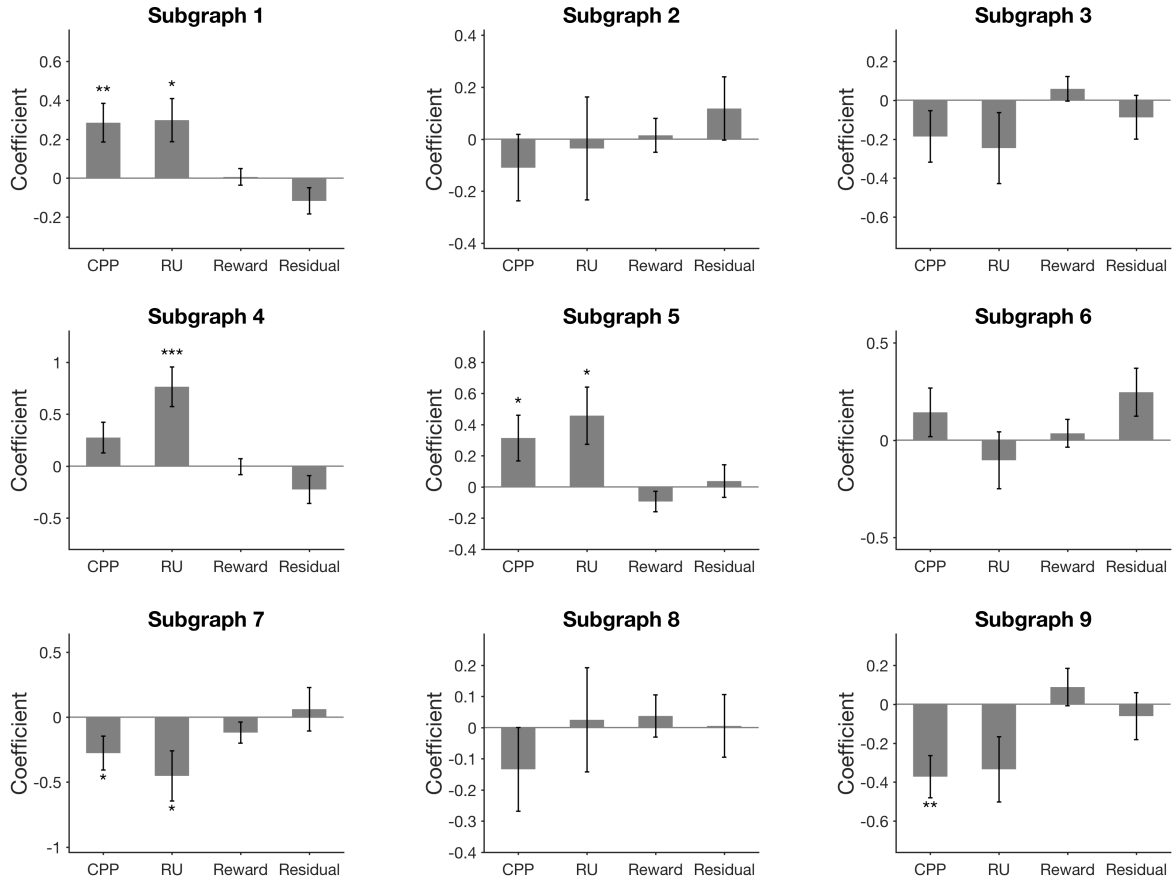
a



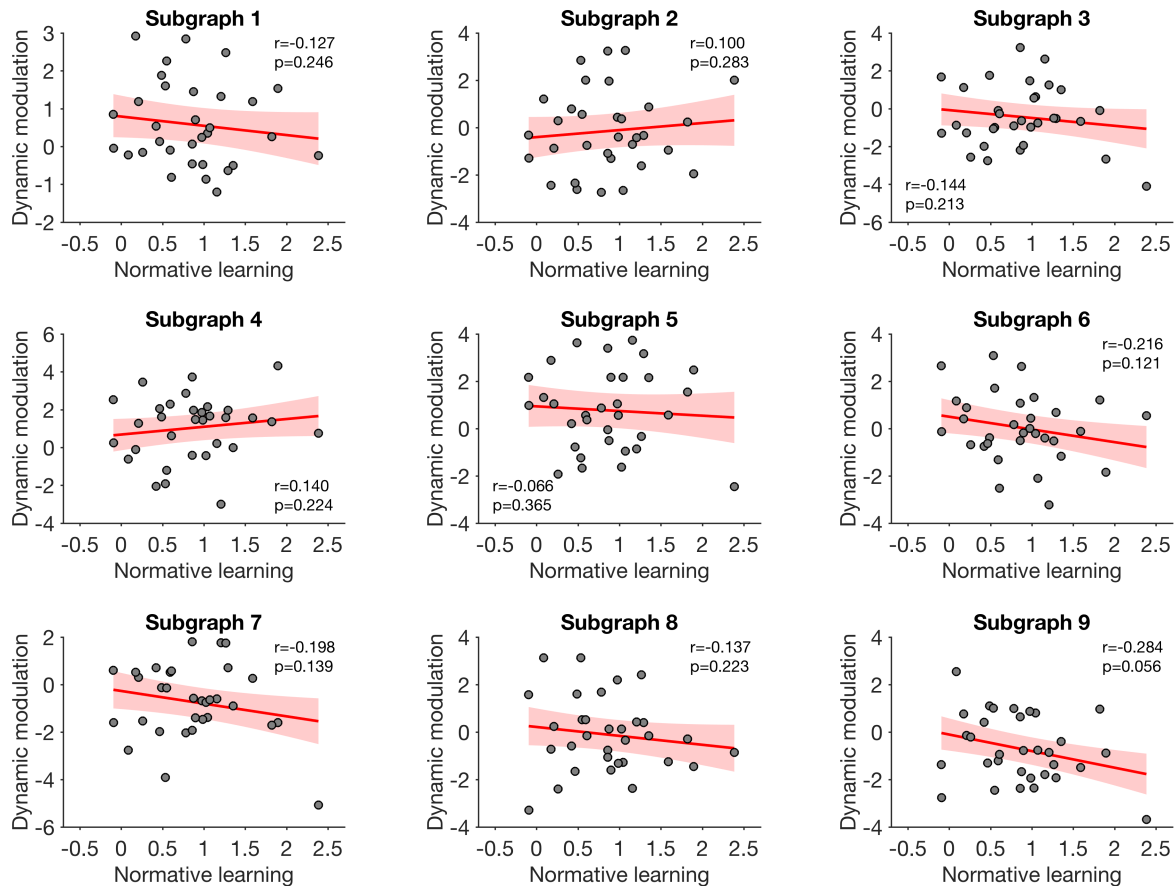
b



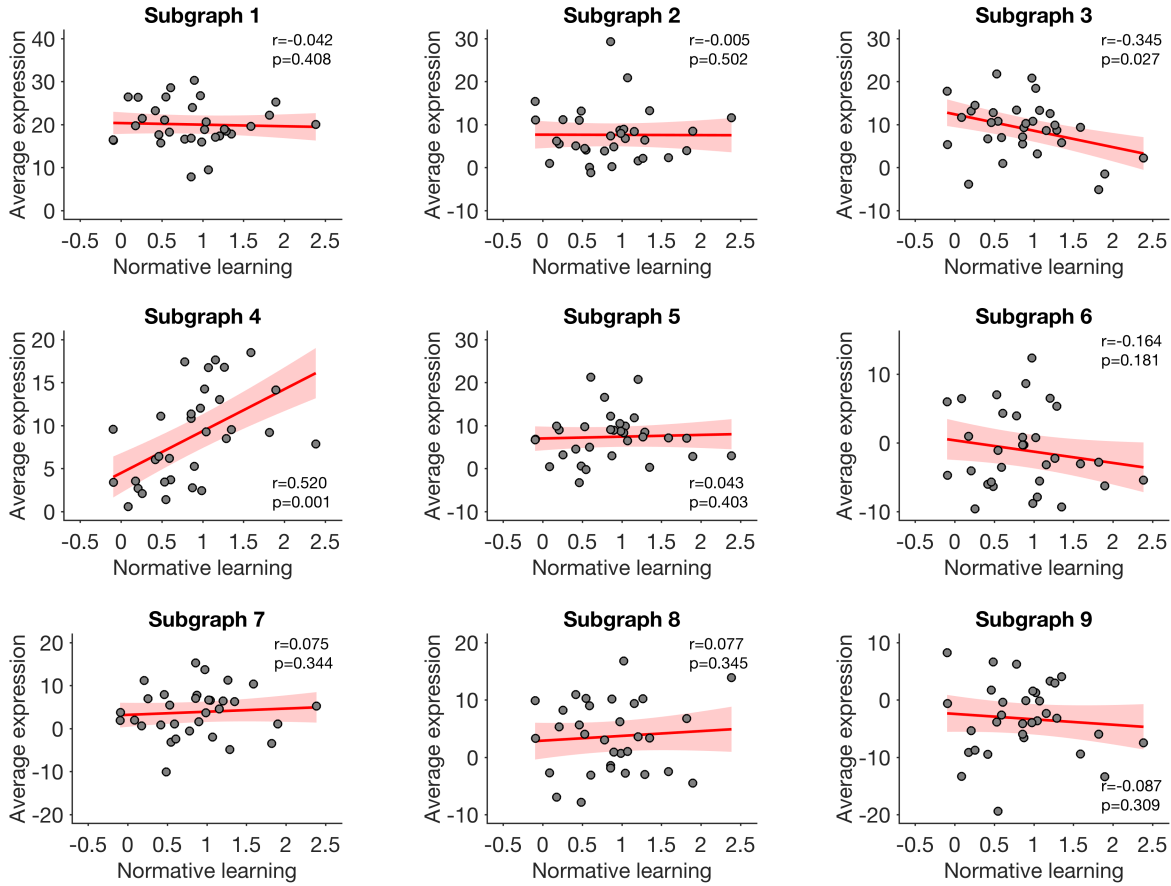
Supplementary Figure 16. Robustness check. Subgraphs identified in predicted BOLD signals from univariate GLMs (including predictors for CPP, RU, reward and residual updating). Subgraphs were identified with a sliding time window of 10 TRs, with 8 TRs overlapping between time windows. (a) Edges between systems in the nine subgraphs identified by NMF, as in Fig. 3b. (b) Summary of the pattern of connectivity in subgraph 4, as in Fig. 4d, showing the within-system strength and between-system strength of each functional system. The 95% confidence interval of each system was estimated by bootstrapping 10,000 times on the edges of that system.



Supplementary Figure 17. Robustness check. Modulation of temporal expression by learning factors in all nine subgraphs identified in predicted BOLD signals (Supplementary Fig. 16). Error bars represent one SEM. (* $p < 0.05$, ** $p < 0.01$, *** $p < .001$)



Supplementary Figure 18. Robustness check. The relationship between individual normative learning and the dynamic modulation of subgraph expression by normative factors in all nine subgraphs identified in predicted BOLD signals (Supplementary Fig. 16). Each point represents one participant. The red line represents the regression line and the shaded area represents the 95% confidence interval.



Supplementary Figure 19. Robustness check. The relationship between individual normative learning and the average subgraph expression in all nine subgraphs identified in predicted BOLD signals (Supplementary Fig. 16). Each point represents one participant. The red line represents the regression line and the shaded area represents the 95% confidence interval.