Supplementary Material

Ventromedial prefrontal-subcortical systems and the generation of affective meaning

Mathieu Roy[1], Daphna Shohamy[1] and Tor D. Wager[2]

[1] Department of Psychology, Columbia University, New York, NY USA

[2] Department of Psychology and Neuroscience and the Institute for Cognitive Science, the University of Colorado, Boulder, CO USA

Corresponding author: Wager, T. (tor.wager@colorado.edu)

This supplement provides additional information on the studies and methods used to summarize ventromedial prefrontal cortical (vmPFC) function across domains.

Study selection

We selected studies from the Neurosynth database [S1], which is available online at http://neurosynth.org. Briefly, the database used for this review contains nearly 4,400 neuroimaging studies from 17 journals, and, for each study, includes both standard Montreal Neurologic Institute (MNI) and Talairach coordinates that describe peak activation locations found in the studies and the full text of the study. Coordinates were extracted using automated text-parsing algorithms and converted to MNI coordinates if necessary. The set of studies reflects a sample of convenience based on the online availability of published reports. A full description of the database, including validation, potential uses, complete methods, and limitations, can be found in [S1].

For the summary reported in this paper, we selected studies from eight functional domains thought to involve the vmPFC. The method for associating studies with each domain was based on the frequency of the use of specified words; for example, for the search query 'pain', studies that used the term 'pain' more than once per 1,000 words would be identified as pain-related, and those that did not would be identified as non-pain related. Although the identification of studies with domains is thus based on the language used in the published papers, the approach yields maps that are very similar to manually coded meta-analyses in several domains (e.g., pain, emotion, and working memory) and can yield useful information. For example, they have high predictive validity in some domains at the level of discriminating the task type performed in individual subjects [S1].

In this review, each domain was identified using a search query that consisted of a combination of search terms. This functionality was not available on the online website at the time of writing, although there are plans to develop flexible search queries for online searches in the future. The list of domains and search queries that identified papers in each domain is listed below. In the list below, |, &, and ~ indicate logical OR, AND, and NOT, respectively. * indicates a wildcard including all words beginning with the stem preceding the star.

Default mode: 'default | resting state | dmn | default mode'

Memory: Average of 'episod*' 'autobiograph*' 'retriev* | recollect*'

Self: '(self | subjective)'

Social cognition/mentalizing: 'empath* | theory.of.mind | tom | mentaliz* | trait | (inference & others)'

Emotion: 'emotion* | mood | valence | arousal | affective'

Reward: '(reward* | monet* | gain | cocaine | eating | reinforc* | incent* | love | joy | (positive & hedonic) | (positive & emotion) | (positive & affect)) &~ (negative & emotion)'

Autonomic/endocrine: 'autonomic* | scr | hr | cortisol | conductance | heart'

Pain: 'pain* | noxious | nocicept*'

For each set of studies, a domain map – a brain image summarizing the activations across studies within each domain – was created, as described below.

Characteristics of the study set

This selection procedure yielded a set of 1,669 unique studies, with numbers in each category ranging from 126 (for the Autonomic domain) to 869 (for the Emotion domain), as described in Table S1. The sets of studies were partially overlapping, as a study could meet the inclusion criteria for more than one category. Several measures were taken to assess the extent of overlap and ensure that 'common' activation in the vmPFC was not an artifact of commonality in the set of studies chosen.

To quantify the overlap in studies, we calculated the Jaccard index for the sets of individual studies used in each pair of maps. The Jaccard index provides a measure of how much overlap there is between sets of studies used in the different domain maps. It provides a measure of percentage overlap between the study sets that contribute to two task-domain maps. It ranges between 0 (no overlap) and 1 (100% overlap), and it is calculated as the intersection of the studies used in a pair of maps divided by the union of all studies used in both maps.

The maximum Jaccard index for any pair across the 8 task domains we used was 0.14, and 85.7% of pairs have a Jaccard index of less than 0.10. Thus, there was relatively little overlap between the sets of images used in each study. An image of the Jaccard matrix for all pairs across the 8 domains is shown in Figure S1. In addition, Table S1 shows the maximum Jaccard index between each domain and any other domain, along with the name of the most overlapping other domain. The table shows that the overlap in study sets that does exist is largely driven by overlap with the 'Emotion' map, indicating that studies from many domains tended to also fit the criteria for the emotion domain, which is unsurprising as emotion is a general category likely to be discussed in multiple literatures.

Table S1.

Map	Studies (N)	Maximum overlap with other maps: Max Jaccard ind	Most overlapping map	Studies unique to this domain
Default mode	148	0.04	Self	93
Memory	296	0.05	Emotion	200
Self	220	0.14	Emotion	39
Social cognition	185	0.13	Emotion	46
Emotion	869	0.14	Self	416
Reward	328	0.12	Emotion	166
Autonomic	126	0.07	Emotion	34
Pain	236	0.09	Emotion	120

Figure S1. Overlap in studies

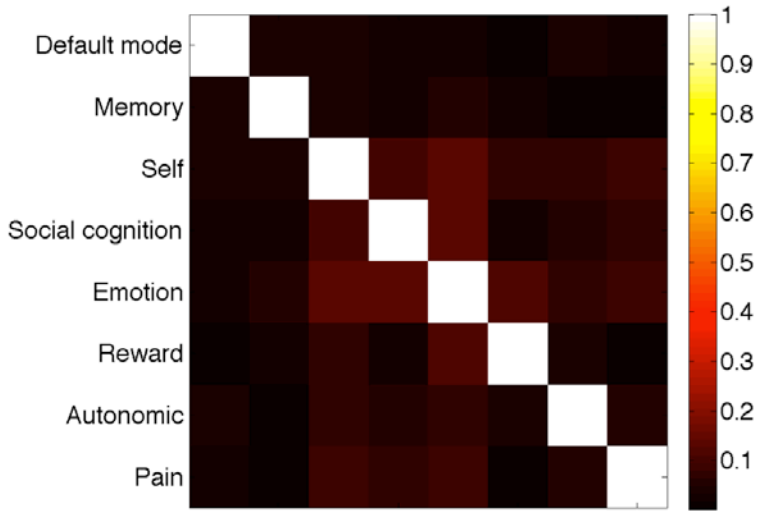

Construction of domain maps and factor analysis (Figure 2)

A domain map for each method was constructed using the 'reverse inference' procedure implemented in Neurosynth [S1]. The domain maps identified regions ('voxels') in which activation of the region significantly predicted whether a study was included in the domain or not. That is, activations in significant regions make it more likely that a study belongs to a domain of interest.

Alhough the analysis is described in detail elsewhere [S1], briefly, the maps were constructed separately for each domain. Each voxel was coded as 'activating' for each study if that study reported an activation coordinate within 10 mm of the voxel. A standard analysis was conducted on the contingency between whether the studies a) activated or did not and b) were identified with the domain or were not. The resulting P-values for each voxel in gray and white matter within a standard MNI brain were saved and thresholded using False Discovery Rate control at $q < .05$. The resulting significant regions are displayed in Figure 2.

To assess the relationships among domain maps and identify potential common underlying components, a matrix consisting of the eight domain maps was subjected to factor analysis with

two factors. Factor analysis decomposes the covariance matrix of the domain maps (an 8 x 8 matrix) in a way similar to principal components and independent components analysis. A principal difference is that in the factor analysis model, each domain is allowed to have its own unique error variance, as well as the variances explained by the common factors. We used factor analysis as implemented in the function factoran.m in Matlab 2010b, with standard varimax rotation and weighted-least squares based calculation of factor scores. Models with more than two factors were not attempted because the model is not identified with more than two factors across eight variables. The goodness of fit was not assessed, as we intended this for descriptive purposes, to illustrate and visualize the two primary components underlying the maps, and not to claim or statistically test the notion that a two-factor solution is adequate.

In the factor analysis, the resulting factor loadings reflect the relationships between the domain maps and the two factors. We visualized these relationships in Figure 2b, in which the axes represent the two factors and the factor loadings for each domain are plotted as points in the factor space. The lines reflect bivariate correlations between pairs of domain maps (independent of the analysis), with domain maps correlated above r = 0.25 connected with lines. The factor scores from the factor analysis are, in this case, maps across brain voxels showing the expression of each factor on each voxel. For each factor map, voxels with loadings in the top 1% of loadings across the brain are visualized on brain slices in Figure 2b.

Assessment of the impact of overlap in studies across domains

Two additional checks were conducted to assess the impact of the overlap in studies on a) involvement of the vmPFC across domains; and b) the factor analysis that identified common components underlying the domain maps. The first check involved constructing domain maps from studies *uniquely* identified in only one of the eight domains. The number of domain-unique studies ranged from 34 (for Autonomic) to 416 (for Emotion), and is listed for each domain in Table S1. The second check involved whitening (decorrelating) the covariance matrix across maps used in the factor analysis using the Jaccard matrix in Figure S1, and examining the similarity of the adjusted covariance matrix to the one used in the main analyses. This allowed us to assess the impact of the overlap in the studies contributing to each domain map on the factor analysis.

Supplementary analysis: Domain maps based on studies uniquely involved in one domain

We constructed domain maps based on studies uniquely involved in only one domain in order to assess whether there was still overlap in the vmPFC across each domain. One caveat is that it is not meaningful to do the analysis to construct the map: Those maps would localize brain areas that respond specifically to studies in one domain that do not meet the criteria for other domains, and so would not be suitable for demonstrating consistent activation in a domain of interest.

Instead, we constructed 'forward inference' maps using the Multilevel Kernel Density Approach, as described in detail in several previous publications [S2]. Briefly, the maps summarize the number of separate studies that activate within 10 mm of each voxel. The maps were thresholded using a Monte Carlo test. The test involved a) randomizing the locations of contiguous activated areas within a gray-matter standard brain mask, b) saving the maximum value across the brain for each of 1,000 iterations, and c) Thresholding based on the 95th percentile of the distribution of maxima to obtain Family-wise Error Rate corrected significance threshold at P < .05 corrected. In addition, we used cluster extent-based correction with a primary threshold of P < .001 and cluster level correction to P < .05 corrected [S2].

The domain maps obtained in this way are shown in Figure S2. The left panels show the individual domain maps, and the right panel shows the overlap across all eight maps. Overlapping regions included the vmPFC, the dorsal anterior cingulate (dACC), and the posterior cingulate. (Though it is hard to see, there are voxels in the Self map in the vmPFC that survive multiple comparisons correction).

There are several reasons that these maps are dispreferred, and the 'reverse inference' maps are presented in the main paper. First, the sample of studies here is considerably smaller. Second, the studies are no longer representative of the set of studies in each domain: for example, Self studies were only included if they used self-related terms frequently but *did not* mention emotion, default mode, etc., which is a potentially biased sample of studies. Third, the reverse inference maps are typically more specific than forward inference maps. For example, dACC is activated in all domains, likely due to the involvement of response and attentional selection in tasks of all types, but it does not appear in the reverse inference maps because it is not *selectively* associated with any of the domains except pain. It is generically activated across these and many other domains, whereas vmPFC activity is selective for each of the eight domains we discuss here. Nonetheless, the forward inference maps converge with the reverse inference maps presented in Figure 2 to suggest that vmPFC is involved in each of the eight domains.

Figure S2.

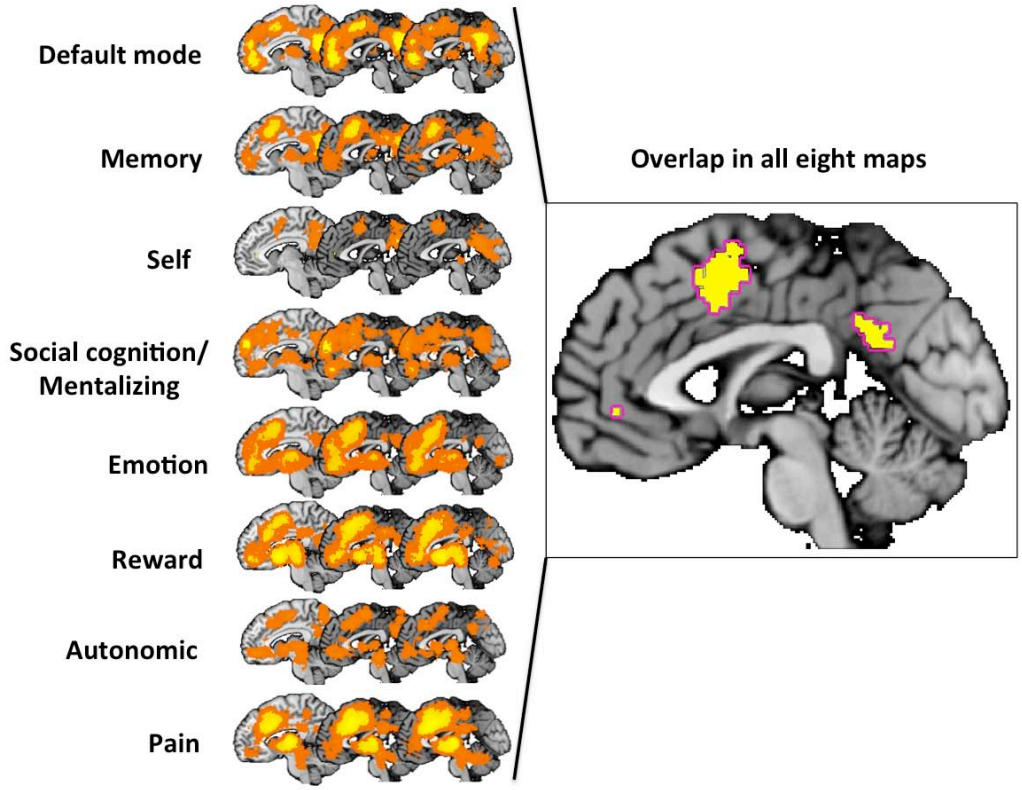

Influence of study overlap on factor analysis

The Jaccard analysis indicated that much of the overlap across studies was due to overlap between the Emotion domain—which unsurprisingly included many studies—and other domains.

Based on this, one would not expect it to drive the clustering of the results we observed, with Reward, Autonomic/Endocrine, and Emotion domains grouped and loading on Factor 1, Default Mode and Memory domains grouped and loading on Factor 2, and Self and Social Cognition domains grouped and loading on both factors (Figure 2b).

To further check this, we whitened the covariance matrix by post-multiplying by the inverse of the Jaccard matrix. The covariances after whitening were correlated with the original covariances r = 0.97, indicating very similar covariances. Repeating the factor analysis on the whitened covariance matrix revealed the same groupings of a) Reward, Autonomic/Endocrine, and Emotion, b) Default Mode and Memory, c) Self and Social Cognition, and d) Pain as separate from the others. The solution was rotated in factor space, as would be expected, but the relationships among the domains were very similar.

Co-activation maps

Co-activation maps were generated to visualize the locations in the MPFC that are typically co-activated with other regions in the brain, and are presented in Figure I in Box 1. Co-activation means that studies that report activations near one brain voxel (within 10 mm) also report activations in another voxel at a rate above what would be expected by chance. Co-activation is a meta-analytic proxy for functional connectivity.

We began by defining 'seed' regions in the dACC, vmPFC, and rdACC based on the peak activity across all studies included in summary. There were distinct areas of high-density activations in dACC and vmPFC in the 1,669 studies used in this review, and seed regions were identified based on the peak density in each of these areas (MNI coordinates [x, y, z] = [2, 14, 48] for dACC and [-2, 42, 0] for vmPFC). The rdACC seed ([0, 30, 32]) was chosen (by T.D.W.) to be an area that lies approximately mid-way between these two seeds and falls within the aMCC subdivision described by Palomero-Gallagher, Vogt and colleagues [S3]. In each case, a 10 mm sphere was drawn around the region of interest's center, and studies activating within 10 mm of an in-region voxel were considered to activate the area.

For each seed region, a standard analysis was conducted on the contingency between whether each study a) activated the region or did not and b) activated within 10 mm of a given voxel elsewhere in the brain or did not. The value was assessed for each voxel, and maps were thresholded at P < .05 Bonferroni corrected ( ). A number of truly co-activated regions are likely missed due to this stringent threshold, particularly small subcortical regions, but the analysis served to illustrate the major patterns of connectivity with the medial prefrontal subregions. The results are shown in Figure Ia in Box 1.

To further investigate co-activation with some of the smaller subcortical nuclei, we defined seed regions in several subcortical areas (Figure Ib in Box 1). Subcortical regions were chosen based on differences in connectivity with different medial prefrontal sub-regions in animal models, but the list of regions is not intended to be comprehensive. In this case, we defined 4 mm spheres around coordinates centered in nuclei localized on the standard single-subject MNI brain (though precise localization of these nuclei is approximate and subject to error based on variation in both anatomy and inter-subject registration procedures). For bilateral areas (amygdala, accumbens, PAG, hypothalamus), the 'seed' was defined as activation within 10 mm of the region in *either* hemisphere. In this case, a more lenient threshold of P < .0001 was used.

Seed coordinates used in Figure Ib are as follows:

Dorsal PAG: [5, -31, -8]; [-3, -31, -8]

Ventral PAG: [5, -26, -9]; [-3, -26, -9]

Dorsal Raphe: [2, -22, -10]

Basolateral amygdala: [-23, -3, -20]; [27, -3, -20]

Nucleus accumbens: [-7, 11, -9]; [12, 11, -9]

Hypothalamus: [-4, 2, -4]; [8, 2, -4]

Positive and negative emotion maps

For the maps illustrated in Figure Id, we used the dataset available from a recent meta-analysis of emotion [S4], which included 234 PET and fMRI studies coded into 374 experimental contrast maps, which formed the unit of observation (a study may report two or more maps corresponding to different experimental tests). From this, we selected 138 studies coded as investigating positive and negative emotional experience [S4].

For each of positive and negative experience, a meta-analysis map was constructed using the MKDA method [S2], which yields a map of the proportion of studies within the category that activate within 10 mm of each voxel, weighted by sample size and analysis quality. In order to compare the positive and negative maps and adjust for overall differences in density across the brain, we standardized each density map by converting to Z-scores across the space of the brain, and subtracted the positive and negative maps. Regions in which the positive vs. negative difference exceeded Z = 2 in either direction were visualized in Figure Id. These maps show differences in the relative density of positive and negative activation peaks, but do not necessarily reflect significant differences in the absolute likelihood of activating each voxel given the emotion category, and they are intended for illustrative and hypothesis-generation purposes.

Difficulties with neuroimaging of vmPFC

Although the current account of vmPFC's functions aims at specifying the processes that should activate it in human studies, several factors complicate the interpretation of activations and deactivations in this area. First, the coarse nature of the BOLD signal cannot differentiate between activations caused by different neuronal populations located in the same voxel. Given that only cortical layers 5/6 project to subcortical structures [S5] and projections synapse on both 'excitatory' neurons and inhibitory interneurons in the subcortex [S6], it is difficult to tell whether or not activation in the vmPFC is having excitatory or inhibitory effects on subcortical structures.

Secondly, the ability to observe activity in the vmPFC is often impaired by the strong BOLD signal dropout and spatial distortion problems affecting this region as a result of its proximity to the sinuses [S7]. Several methodological measures can be taken to reduce the impact of such problems, but these are not systematically applied, and are employed more often in some domains than others. Thus, the lack of vmPFC activation, or the precise location of such activations, should be interpreted with caution, especially when comparing studies performed with different acquisition parameters. In addition to false negatives in vmPFC due to signal

dropout, typical echo-planar BOLD acquisition schemes produce systematic spatial distortions that produce an artifactual dorsal shift in vmPFC activation. These artifacts contribute to both the spatial variability across studies (Compare Figure 6 in [S8] and Figure 6 in [S9], which studied the same social threat task with echo-planar and spiral imaging techniques) and could cause activations in vmPFC to be localized to the rdACC or even dACC.

Finally, activations in the vmPFC can be offset by other factors, such as deactivations associated with the performance of effortful external tasks [S10]. Such deactivations may reflect antagonism between the processes involved in directed attention and internal self-reflection, or they may reflect separate effects in the same gross anatomical region. Thus, tasks in which meaning-generation is effortful, such as voluntary reappraisal of negative emotion, could show little vmPFC activity if meaning-generation demands a great deal of directed attention or effort (which reappraisal does; [S11,S12]). In such cases, task-related decreases could mask meaning generation-related increases. Consequently, lack of vmPFC activation in certain experiments involving the construction of meaning should be interpreted with caution.

References

S1 Yarkoni, T. *et al.* (2011) Large-scale automated synthesis of human functional neuroimaging data. *Nat. Methods* 8, 665–670

S2 Wager, T. D. *et al.* (2009) Evaluating the consistency and specificity of neuroimaging data using meta-analysis. *NeuroImage* 45, S210–S221

S3 Palomero-Gallagher, N. *et al.* (2009) Receptor architecture of human cingulate cortex: Evaluation of the four-region neurobiological model. *Hum. Brain Mapp.* 30, 2336–2355

S4 Lindquist, K. *et al.* The brain basis of emotion: a meta-analytic review. *Behav. Brain Sci.*

S5 Vertes, R. P. (2003) Differential projections of the infralimbic and prelimbic cortex in the rat. *Synapse* 51, 32–58

S6 Pare, D. (2004) New Vistas on Amygdala Networks in Conditioned Fear. *J. Neurophysiol*. 92, 1–9

S7 Huettel, S. A. *et al.* (2003) *Functionnal Magnetic Resonance Imaging*, Sinauer Associates, Inc.

S8 Wager, T. D. *et al.* (2009) Brain mediators of cardiovascular responses to social threat. *NeuroImage* 47, 821–835

S9 Wager, T. D. *et al.* (2009) Brain mediators of cardiovascular responses to social threat, Part II: Prefrontal-subcortical pathways and relationship with anxiety. *NeuroImage* 47, 836–851

S10 Raichle, M. *et al.* (2001), A default mode of brain function. in *Proc. Natl. Acad. Sci. U. S. A.*, 98, pp. 676–682

S11 Urry, H. L. *et al.* (2009) Individual differences in some (but not all) medial prefrontal regions reflect cognitive demand while regulating unpleasant emotion. *NeuroImage* 47, 852–863

S12 van Reekum, C. M. *et al.* (2007) Gaze fixations predict brain activation during the voluntary regulation of picture-induced negative affect. *NeuroImage* 36, 1041–1055