

The dorsal anterior cingulate cortex is selective for pain: Results from large-scale reverse inference

Matthew D. Lieberman¹ and Naomi I. Eisenberger

Department of Psychology, University of California, Los Angeles, CA 90095-1563

Edited by Richard Ivry, University of California, Berkeley, CA, and accepted by the Editorial Board October 26, 2015 (received for review July 30, 2015)

Dorsal anterior cingulate cortex (dACC) activation is commonly observed in studies of pain, executive control, conflict monitoring, and salience processing, making it difficult to interpret the dACC's specific psychological function. Using Neurosynth, an automated brainmapping database [of over 10,000 functional MRI (fMRI) studies], we performed quantitative reverse inference analyses to explore the best general psychological account of the dACC function P(Ψ process|dACC activity). Results clearly indicated that the best psychological description of dACC function was related to pain processing—not executive, conflict, or salience processing. We conclude by considering that physical pain may be an instance of a broader class of survival-relevant goals monitored by the dACC, in contrast to more arbitrary temporary goals, which may be monitored by the supplementary motor area.

dACC | pain | reverse inference | Neurosynth

Of all the regions in the brain that have received heavy study, there may be no region with less consensus about its specific function than the dorsal anterior cingulate cortex (dACC). This lack of consensus is not for lack of interest. The cingulate cortex is a mammalian-specific region (1) that hugs the entire length of the corpus callosum, and the dACC (Fig. 1A) constitutes a section of the cingulate that sits above the corpus callosum, with the cingulate sulcus as its dorsal boundary [in Montreal Neurological Institute (MNI) coordinate space, $0 \leq y \leq 30$ defines its approximate anterior/posterior boundaries]. Over the past few decades, there have been various lines of inquiry that have all claimed that the dACC plays a central role in in any one of several processes of interest. These domains include executive processing, working memory, inhibitory control, conflict monitoring, pain, emotion, consciousness, and, most recently, salience (2–12). Using Neurosynth (13), an automated brainmapping database, we performed quantitative reverse inference analyses to explore the best general psychological account of dACC function.

Perhaps the biggest difficulty with prior claims about the psychological function of the dACC is that these claims have relied on forward inference data to draw conclusions about the functions of particular regions. Forward inference, in this context, refers to the probability that a study or task that invokes a particular process will reliably produce dACC activity [e.g., the probability of dACC activity, given a particular psychological process: $P(\text{dACC activity}|\Psi \text{ process})$]. Forward inference data can be generated in traditional neuroimaging meta-analyses. Strong forward inference to dACC has been shown in each domain mentioned above, which may be why some researchers have suggested that activity observed in the dACC cannot be reliably linked to any psychological process because, on the surface, it seems that many different psychological processes activate this region.

Although forward inference is important and typically identifies regions that are plausible candidates to contribute to a particular psychological process, it is a logical error (i.e., affirming the consequent) to draw psychological conclusions about regions present in forward inference analyses (14). Just because psychological process A reliably produces activity in region X does not mean that activity in region X in a new dataset indicates that psychological process A was invoked.

To identify the psychological contributions of a region, reverse inference analyses, rather than forward inference analyses, are needed. Although reverse inference has been used as a derogatory term to characterize the error of using forward inference data to draw conclusions about a region's function (“that's just bad reverse inference”), appropriate reverse inference methods exist and allow the desired inferences to be drawn. Reverse inference, in the current context, refers to the probability that dACC activity can be attributed to a particular psychological process [i.e., the probability of a given psychological process, given activity in the dACC: $P(\Psi \text{ process}|\text{dACC activity})$].

Until recently, large-scale reverse inference analyses were quite difficult to carry out. Even popular multivoxel pattern analysis (MVPA) approaches are limited to reverse inferences only about the specific tasks included in the study. For instance, an MVPA study of pain vs. no pain trials would likely show that dACC has positive predictive value for predicting future pain vs. no pain trials (15). However, such a study would not inform whether a pain account of dACC is better than a working memory account of dACC. Furthermore, an MVPA study of pain and working memory trials would still not inform whether dACC is selective for pain or working memory relative to all of the untested other accounts (e.g., conflict, salience, etc.).

In contrast to MVPA, Neurosynth and other related large-scale neuroimaging databases now offer the opportunity to perform comprehensive reverse inference analyses that include virtually every psychological process that has been attributed to dACC. Neurosynth uses terms that appear frequently in articles as a proxy for the psychological process under investigation in that study. Using a naive Bayesian classifier, Neurosynth is able to identify which psychological processes, from among hundreds simultaneously considered, were likely to have been invoked when activity in a particular brain region was present across 10,903 studies in the database (as of June

Significance

No neural region has been associated with more conflicting accounts of its function than the dorsal anterior cingulate cortex (dACC), with claims that it contributes to executive processing, conflict monitoring, pain, and salience. However, these claims are based on forward inference analysis, which is the wrong tool for making such claims. Using Neurosynth, an automated brainmapping database, we performed reverse inference analyses to explore the best psychological account of dACC function. Although forward inference analyses reproduced the findings that many processes activate the dACC, reverse inference analyses demonstrated that the dACC is selective for pain and that pain-related terms were the single best reverse inference for this region. This finding has implications for our understanding of pain and distress-related psychological disorders.

Author contributions: M.D.L. analyzed data; and M.D.L. and N.I.E. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. R.I. is a guest editor invited by the Editorial Board.

¹To whom correspondence should be addressed. Email: lieber@ucla.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1515083112/-DCSupplemental.




Fig. 1. The dorsal anterior cingulate cortex (dACC). (A) The anatomical outlines of the dorsal anterior cingulate cortex (dACC), rostral anterior cingulate cortex (rACC), supplementary motor area (SMA), and presupplementary motor area (pre-SMA). (B) The Neurosynth forward inference map for the term “dACC” showing its peak effect (noted by the brightest part of the activation cluster) in the SMA/pre-SMA (not in the dACC). (C) The Neurosynth forward inference map for the term “anterior cingulate” also showing its peak effect in the SMA/pre-SMA. Maps use Neurosynth’s standard false discovery rate (FDR) criterion of $P < 0.01$.

2015). A number of papers have now used Neurosynth and related tools to conduct reverse inference analyses that were largely impossible a decade ago (16–20). Here, we use Neurosynth to investigate the psychological processes specific to the dACC by (i) examining the specific location in the brain most associated with the term “dACC,” (ii) conducting forward inference analyses to replicate prior findings that many different psychological processes activate the dACC, and (iii) conducting reverse inference analyses to identify the psychological processes specific to dACC activation.

Step 1: Locating the dACC

Before using Neurosynth to recapitulate prior forward inference analyses generated in support of different claims about dACC functionality, we first examined the specific neural location associated with the term “dACC” in the Neurosynth database. Prior work has noted the close alignment between anatomically defined and Neurosynth-derived boundaries (13). Because we were interested in what regions were active when the term “dACC” was used [P(dACC activity|“dACC” term used)], we conducted a forward inference analysis of the term “dACC.” What was striking in this dACC analysis is that the neuroimaging literature may have mislabeled the location of the dACC in many studies, and this mislabeling may have played a part in the current misunderstanding regarding dACC functionality.

As noted earlier, the cingulate sulcus is the dorsal boundary of the dACC. Above this sulcus are the supplementary and presupplementary motor areas (SMA and pre-SMA), which cover

roughly the same area as the medial aspects of Brodmann areas 6 and 8 (see the green area of Fig. 1A) and are contained within the superior frontal gyrus. When “dACC” is entered as a term into a Neurosynth forward inference analysis (Fig. 1B), there is substantial activity present in the anatomically defined dACC region; however, there is also substantial activity present in the SMA/pre-SMA region. Moreover, the location with the highest Z-score in this analysis is actually in SMA, not dACC. The same is true if the term “anterior cingulate” is used (Fig. 1C). We tested several other anatomical terms including “amygdala,” “hippocampus,” “posterior cingulate,” “basal ganglia,” “thalamus,” “supplementary motor,” and “pre sma.” In each of these regions, the location with the highest Z-score was within the expected anatomical boundaries. Only within the dACC did we find this distortion. These results indicate that studies focused on the dACC are more likely to be reporting SMA/pre-SMA activations than dACC activations. For more on anatomical considerations, see [Supporting Information](#).

These findings suggest that some of the disagreement over the function of the dACC may actually apply to the SMA/pre-SMA, rather than the dACC. In fact, a previous paper reporting that a reverse inference analysis for dACC was not selective for pain, emotion, or working memory (see figure 3 in ref. 13) seems to have used coordinates for the dACC that are in fact in the SMA/pre-SMA (MNI coordinates 2, 8, 50), not in the dACC.

Step 2: Forward Inference in the dACC

Before moving to reverse inference analyses, we started with forward inference analyses. This approach was taken to assess whether Neurosynth would recapitulate the standard problem with dACC activation interpretation: namely, that a variety of different psychological processes reliably activate this region. We first examined forward inference maps for many of the psychological terms that have been associated with dACC activity. These terms were in the categories of pain (“pain,” “painful,” “noxious”), executive control (“executive,” “working memory,” “effort,” “cognitive control,” “cognitive,” “control”), conflict processing (“conflict,” “error,” “inhibition,” “stop signal,” “Stroop,” “motor”), and salience (“salience,” “detection,” “task relevant,” “auditory,” “tactile,” “visual”). See [Table S1](#) for the number of studies associated with each term in this article.

As can be seen in Fig. 2, each term in each of these four categories produced at least some reliable activity in the dACC although, even here, it is noticeable that, for many of the terms, the lion’s share of the activity is dACC-adjacent in the SMA or pre-SMA. Specifically, the forward inference maps for “working memory,” “effort,” “stop signal,” “Stroop,” “salience,” “auditory,” and “tactile” each have a relatively modest footprint within the dACC proper, compared with the other terms of interest. Nevertheless, in general, these forward inference maps support the general concern that many have raised—that various psychological tasks and processes activate the dACC.

Step 3: Reverse Inference in the dACC

We next turned to reverse inference maps for the same set of terms. Importantly, these results looked very different (Fig. 3). Although there was strong evidence of forward inference from a variety of psychological processes and tasks to dACC, there was very little support for reverse inference from dACC activity to almost all of the same processes and tasks. This finding is seemingly consistent with the received wisdom that, if the dACC is activated by everything, then it must be selective for nothing. However, there was one clear exception. The three pain terms (“pain,” “painful,” “noxious”) all showed broad coverage across the dACC in the reverse inference analyses. In other words, whereas psychological processes and tasks related to pain, executive processes, conflict, and salience all reliably activate the dACC, the only psychological phenomenon that can be reliably inferred given the presence of dACC activity is pain. It is noteworthy that the pattern of effects is

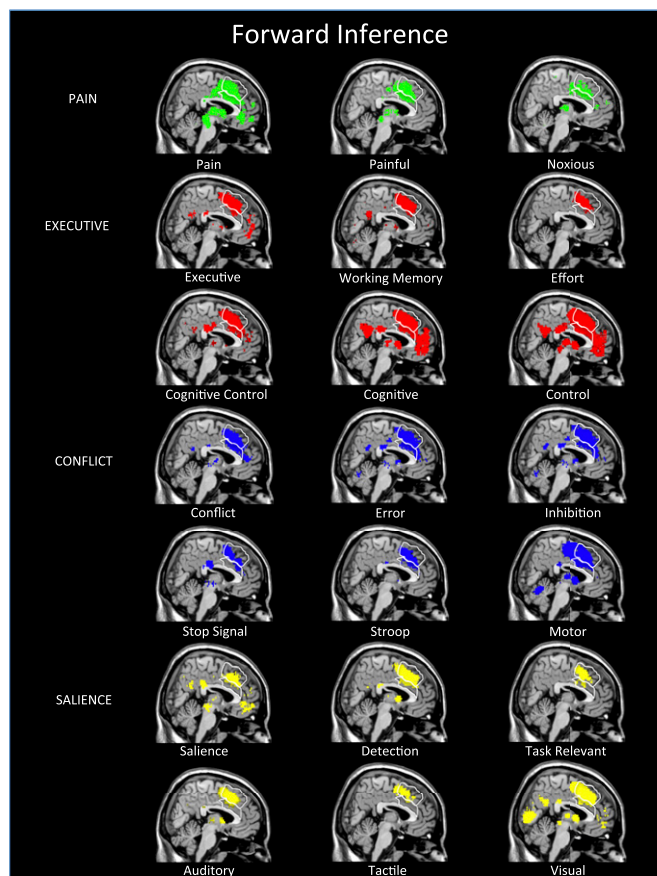


Fig. 2. Neurosynth forward inference maps link several processes to the dACC. Forward inference maps reflect the probability of a region being present when a particular psychological process or task is invoked [$P(\text{dACC activity}|\Psi \text{ process})$]. Shown here are forward inference maps for several terms in the categories of pain, executive, conflict, and salience processes that have all been linked to the dACC in prior work. Outlines show the anatomical boundaries for the dACC (bottom) and the SMA/pre-SMA (top). Maps use Neurosynth's standard FDR criterion of $P < 0.01$.

similar for the anterior insula (AI) as well, a region with similar interpretational concerns (Fig. 4).

Our next goal was to quantify the strength of evidence for different processes being the psychological interpretation for dACC activity and how the evidence for different psychological processes compared with one another. We wanted to explore this issue in an unbiased way across the dACC that would allow each psychological domain to show where there is more or less support for it as an appropriate psychological interpretation. To perform this analysis, we extracted reverse inference statistics (Z-scores and posterior probabilities) across eight foci in the dACC for the terms “pain” (410 studies), “executive” (531 studies), “conflict” (246 studies), and “salience” (222 studies). The foci were equally spaced out across the midline portion of the dACC (see Fig. 5 for coordinates).

We plotted the posterior probabilities at each location for each of the four terms, as well as an average for each psychological term across the eight foci in the dACC (Fig. 5). Because Z-scores are less likely to be inflated from smaller sample sizes than the posterior probabilities, our statistical analyses were all carried out on the Z-scores associated with each posterior probability (21). A posterior probability is akin to an effect size, although not a directly interpretable one, because the Bayesian prior for each term was normed to 0.50. Thus, a posterior probability of 0.82 is likely a significantly larger effect size than another of 0.56; however, due to norming, one cannot say that the 0.82 implies that there is an

82% chance that an activation came from a study with a particular psychological term. In contrast, the associated Z-scores represent the accumulation of evidence that the reverse inference to a psychological process is reliable. It should be noted that the bar graphs in Fig. 5 look substantively the same if Z-scores are plotted instead of posterior probabilities. In addition, we used a threshold of $Z > 3.1$, $P < 0.001$ as our threshold for indicating significance. This threshold was chosen instead of Neurosynth's more strict false discovery rate (FDR) correction to maximize the opportunity for multiple psychological terms to “claim” the dACC. Indeed, when using Neurosynth's stricter FDR correction, no terms other than “pain” were significant at any of these locations.

Several things were noteworthy in these analyses. First, there was strong evidence that dACC activity in seven out of eight foci (all but coordinates 0, 26, 16) could be attributed to pain by quantitative reverse inference. In contrast, there was evidence that dACC activity could be attributed to conflict (in addition to pain) in one out of eight foci (coordinates 0, 18, 30). None of the eight foci within the dACC indicated that activity could be reliably attributed to executive or salience processes.

It is also critical to compare the Z-score for pain with the Z-scores for the other terms. It is possible that the Z-scores for two terms are not significantly different from one another despite one reaching significance and the other not reaching significance on its own.

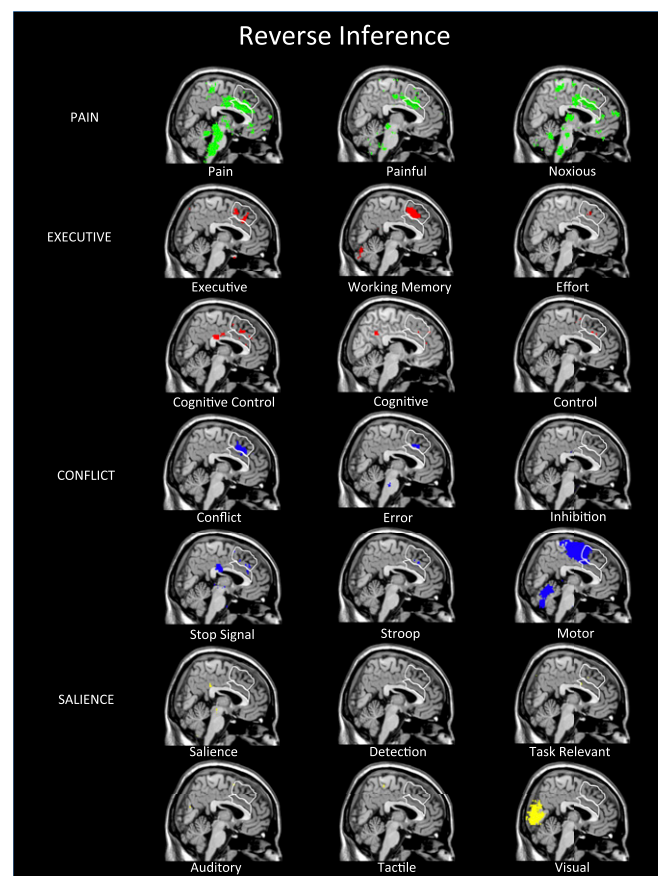


Fig. 3. Neurosynth reverse inference maps indicate that the dACC is selective for pain. Reverse inference maps reflect the probability that a particular psychological process or task was invoked given the presence of activity in a particular region [$P(\Psi \text{ process}|\text{dACC activity})$]. Reverse inference maps for executive, conflict, and salience processes show almost no evidence that dACC activity can be explained in terms of those processes. In contrast, these maps show clear evidence that the dACC can be reliably linked to pain processes. Maps use Neurosynth's standard FDR criterion of $P < 0.01$.

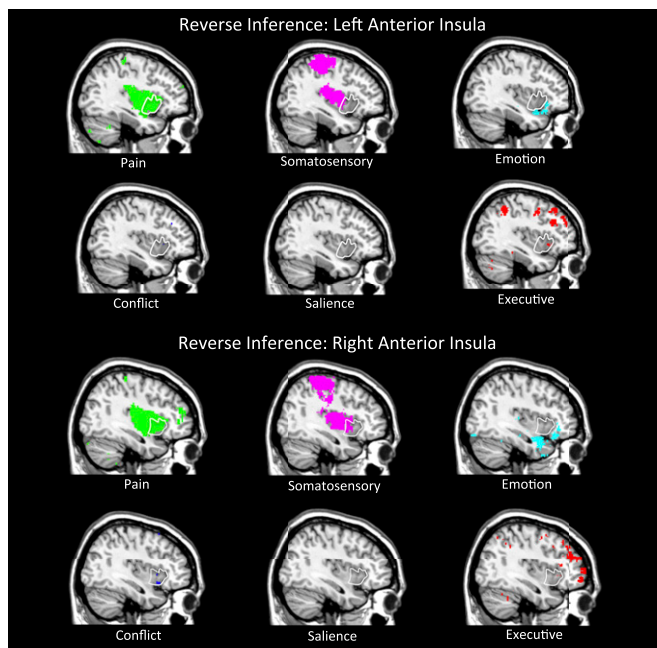


Fig. 4. Neurosynth reverse inference maps indicate the anterior insula (AI) is selective for pain. Reverse inference maps covering bilateral AI are shown for the terms “pain,” “somatosensory,” “emotion,” “conflict,” “salience,” and “executive.” Only the terms “pain” and “somatosensory” show substantial coverage of the AI. However, in both hemispheres, most of the AI shows exclusive effects for pain. Outlines show the anatomical boundaries for the AI. Maps use Neurosynth’s standard FDR criterion of $P < 0.01$.

However, that was not the case here. In each of the seven foci for which there was significant evidence that dACC activity could be attributed to pain, the Z-score for pain was significantly greater than the Z-score for each other term at that foci [$Z_s > 3.1$, $P < 0.001$ using the formula $(Z_1 - Z_2)/\sqrt{2}$] (22). In addition, the maximum pain Z-score for the dACC (11.98) at coordinates 0, 20, 28 was significantly higher than the maximum Z-score anywhere in the brain for the terms “executive” (6.67), “conflict” (5.46), and “salience” (6.09) ($P < 0.001$). In other words, this analysis suggests that there is stronger evidence that dACC activity can be attributed specifically to pain than there is evidence that executive, conflict, and salience processes can be attributed to any specific region of the brain at all. When the data from the eight foci were averaged together, the same pattern emerged, with only pain showing significant support as a target for dACC reverse inference, and this effect was significantly stronger for pain than for the other terms.

Beyond the specific terms we selected for analyses, we also identified which psychological term was associated with the highest Z-score for each of the eight dACC locations across all of the psychological terms in the Neurosynth database. Despite the fact that there are several hundred psychological terms in the Neurosynth database, “pain” was the top term for six out of eight locations in the dACC. For one of the other two locations, “pain” was the second term after “clinically” (coordinates 0, 2, 32). For the last location, “pain” was not in the top 20 terms; however, the top term was “chronic pain” (coordinates 0, 26, 16). Finally, it is notable that the maximum Z-score (11.98) for pain anywhere in the dACC reverse inference maps is in the same range as the maximum Z-score for several other less controversial anatomy-to-function mappings, such as “visual” (13.30) in the visual cortex, “memory” (10.41) in the hippocampus, or “mentalizing” (11.41) in the dorsomedial prefrontal cortex.

The conclusion from the Neurosynth reverse inference maps is unequivocal: The dACC is involved in pain processing. When only forward inference data were available, it was reasonable to

make the claim that perhaps dACC was not involved in pain per se, but that pain processing could be reduced to the dACC’s “real” function, such as executive processes, conflict detection, or salience responses to painful stimuli. The reverse inference maps do not support any of these accounts that attempt to reduce pain to more generic cognitive processes. As seen in Fig. 3, a whole slew of such terms show little or no evidence of being likely candidates to explain the psychological bases of dACC activity. Although some of these functions might be instantiated in the SMA or pre-SMA, they do not seem to be reliably related to dACC activity in quantitative reverse inference analyses. Of the 21 terms we began with, only “pain,” “painful,” and “noxious” were reliable reverse inference candidates for most of the dACC (for more on affective versus sensory aspects of pain, see [Supporting Information](#)).

Neural Alarm Account of dACC

Over a decade ago, we (23) suggested that the dACC was involved in both affective (e.g., distress) and cognitive (e.g., conflict detection) processes, in contrast to the prevailing view at that time that it was primarily involved in cognitive processes (24). More recently, meta-analyses have confirmed that studies of negative affect (in addition to pain and cognitive control) reliably activate

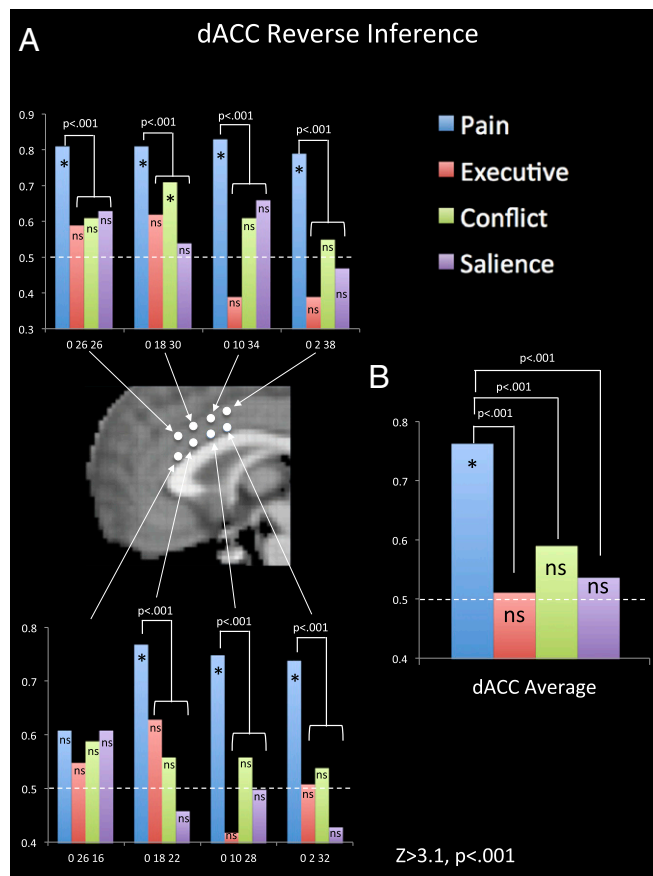


Fig. 5. Comparison of reverse inference effects throughout the dACC. (A) Plotted posterior probabilities from Neurosynth reverse inference maps for pain, executive, conflict, and salience processes at eight foci on the midline of dACC. MNI coordinates are listed on the bar graphs. The dashed lines at 0.50 indicate the null hypothesis (i.e., no reverse inference evidence). All starred bars (*) had Z-scores of ≥ 3.1 , $P < 0.001$, and “ns” indicates not significant at this threshold. For each location where “pain” was a reliable reverse inference term, “pain” was also a significantly stronger reverse inference term than all other terms ($P < 0.001$). (B) The average posterior probabilities across the eight foci for pain, executive, conflict, and salience.

the dACC (8, 9). In our earlier account, we suggested that the dACC served as a neural alarm notifying us when a goal-related conflict occurs that requires our attention. We posited both a cognitive function (i.e., detecting the conflict) and an affective function (i.e., sounding the alarm) analogous to any kind of mechanical alarm. This account was originally intended as a way to integrate pain and conflict accounts of the dACC, rather than simply explaining the former in terms of the latter. Supporting this view, we observed that in a Stop Signal task, making errors produced dACC activity; however, the extent that these errors distressed participants modulated the dACC over and above the effect of the occurrence of errors (25) (see also refs. 26 and 27).

Neurosynth reverse inference maps suggest that, whereas the sounding of the alarm may be specific to the dACC, the conflict detection itself is more strongly associated with the SMA, dorsal to the dACC. We used Neurosynth to produce reverse inference maps of several emotion-related terms that might be related to the alarm-sounding function, including “emotion,” “negative affect,” “negative,” “distress,” “fear,” “anger,” “anxiety,” and “depression,” along with “reward,” another affective term that has also been linked to the dACC in previous studies (Fig. S1). Distress-related emotions (“negative affect,” “distress,” “fear”) were each linked to a dACC cluster, albeit much smaller than the one associated with “pain,” whereas other emotion terms generally were not. “Reward” was the one exception; however, reward was associated with the most anterior portion of the dACC, with little overlap in the portion of dACC linked to “pain” in the reverse inference maps. In contrast, Fig. 3 shows clearly that reverse inference maps of conflict detection-related terms produce nearly all of their effects in the SMA/pre-SMA, rather than in the dACC proper. Although these findings are still consistent with a two-step neural alarm process, they are less supportive of both steps occurring in the dACC.

Arbitrary vs. Survival Goal Conflicts

We suggest here a new account of dACC function relative to the SMA region that is just above it. To be clear, this new account is the most speculative aspect of the current article, but is a new proposal worthy of further investigation. First, it must be noted that the dACC is phylogenetically older than the SMA/pre-SMA with a simpler laminar structure (28, 29). Second, in reverse inference maps both dACC and SMA/pre-SMA activity point to psychological responses to goal conflicts; however, they seem to respond to qualitatively different forms of conflict. Avoiding pain is an ever present survival-relevant goal. In contrast, most conflict detection and error studies use arbitrary temporary goals (“Don’t press the button when you see an X”).

This account suggests that the dACC may have evolved earlier to respond to enduring survival-relevant goals that support hardwired survival functions (30, 31). In contrast, the SMA/pre-SMA may have evolved to detect conflicts interfering with the kind of abstract temporary goals generated in the lateral prefrontal cortex. If this distinction successfully characterizes dACC vs. SMA/pre-SMA functionality, then it would be expected that, when other hardwired survival-relevant goals are threatened, including those that do not involve nociception, the dACC would be activated.

We have argued that social rejection represents a threat to the survival-relevant goal of social connection (32, 33). Recently a meta-analysis (34) was published on 33 studies of social rejection and found that the strongest reliable response ($P_{FDR} < 0.001$) was at the coordinates (8, 24, 24) in the dACC in a region for which the strongest term in the Neurosynth reverse inference map is “pain.”

Beyond social rejection, there are a number of other survival-relevant goal conflicts relevant to our hypothesis such as hunger, thirst, and breathlessness. There are only a handful of studies that have examined these processes to date (35–41), and most are PET studies that can more easily compare different enduring states than functional MRI (fMRI). As shown in Fig. 6, all of these studies produce activity in the dACC. Although not reverse

inference-based, because “hunger,” “thirst,” and “breathlessness” are not in the Neurosynth database of terms, these results are consistent with the proposed division of labor between the dACC and SMA. The dACC seems to play more of a role in enduring survival-relevant goal conflicts whereas the SMA plays more of a role in flexible temporary goal conflicts.

Limitations

Neurosynth is not a perfect tool, and its limitations have been written about previously (13). As an automated tool, it does not perform content analyses of how terms are being used in a paper. Both activations and deactivations that appear in tables will be treated the same way. Additionally, Neurosynth is affected by the same confirmation biases that affect other tools. If researchers expect certain regions to be active for certain tasks, they are more likely to report activations in those regions rather than others. Finally, the reverse inference is linguistic, focused on the terms used across articles rather than on task trial types of specific psychological states. Nevertheless, the creators of Neurosynth provide several demonstrations (13) that these limitations do not prevent Neurosynth from providing robust quantitative reverse inference data consistent with other databases and methods of analysis. Although no tool is perfect for every job, Neurosynth is a powerful tool that can help answer previously hard-to-address questions.

Another potential concern about the current findings is that different terms have different base rates within Neurosynth. It is possible that terms that occur more frequently, like “pain,” might naturally produce stronger reverse inference effects than less frequent terms. This concern is addressed in two ways. First, the current analyses included a variety of terms that included both more or fewer studies than the term “pain” and no frequency-based gradient of dACC effects is observable. Second, Neurosynth explicitly controls for frequency by setting the Bayesian prior for every term in the database to 0.50. This procedure prevents high frequency terms from overwhelming rare terms in reverse inference analyses (see *Supplementary Information* in ref. 13 for a full discussion).

A related issue concerns the specificity of different terms in our analyses. It is possible that more tightly defined constructs produce reverse inference effects more readily than more loosely defined constructs, which may reflect a collection of distinct processes. Although in principle this concern could affect the outcome of our reverse inference analyses, two points argue against it. First, across the different terms analyzed in this study, both loosely defined (e.g., “executive”) and tightly defined (e.g., “conflict detection”) constructs failed to produce dACC reverse inference

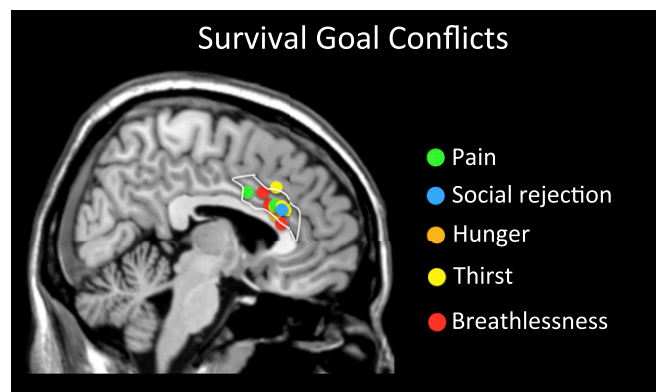


Fig. 6. Survival-relevant goal conflicts in the dACC. Several survival goal conflicts produce activations in the dACC. Pain foci were derived from the two strongest reverse inference effects from Neurosynth. Social rejection activation comes from a meta-analysis (34). Hunger (35), breathlessness (36–38), and thirst (39–41) foci all come from individual studies.

effects. Second, although loosely defined constructs, such as “executive,” “motor,” and “visual,” do not produce dACC responses, each produces strong reverse inference effects in other expected regions of the brain. This finding argues against the notion that these terms were being biased against in our dACC analyses.

A final issue focuses on individual variability in neuroanatomy. Some individuals have an additional paracingulate sulcus, which would extend the dACC more dorsally (42). Given that Neurosynth does not take individual anatomical variability into account, it is possible that, for some individuals, the dACC does extend to the areas where some terms other than “pain” are showing reverse inference effects. Even in this case, there would still be strong evidence of a dorso-ventral gradient distinguishing pain effects spanning all of the dACC and some nonpain effects at the dorsal-most aspect of the dACC (see [Supporting Information](#) for further discussion of the paracingulate sulcus).

Conclusions

Cognitive neuroscience has long lived in a world of forward inference, with many tasks activating the dACC. Neurosynth’s

forward and reverse inference analyses clearly show several things: (i) Articles using the term “dACC” may be mistakenly labeling the SMA/pre-SMA as the dACC; (ii) numerous tasks reliably activate the dACC; but (iii) the best interpretation of dACC activity is in terms of pain processes; and (iv) the evidence for this pain-based account is significantly stronger than for alternative accounts, such as executive, conflict, and salience processes. Based on available evidence, the clearest account of dACC function is that it is selectively involved in pain-related processes.

Materials and Methods

Neurosynth (www.neurosynth.org) was used to conduct forward and reverse inference analyses on 10,903 neuroimaging studies that were available as of June 2015. Analytic details are included in the main text; the numbers of studies associated with each search term are listed in [Table S1](#). Informed consent was unnecessary, and institutional approval was not needed in this study because no new data were collected.

ACKNOWLEDGMENTS. We acknowledge J. Rissman and E. Berkman for assistance in interpreting Neurosynth results and T. Yarkoni with providing Neurosynth assistance.

- MacLean PD (1985) Brain evolution relating to family, play, and the separation call. *Arch Gen Psychiatry* 42(4):405–417.
- Vogt BA, Finch DM, Olson CR (1992) Functional heterogeneity in cingulate cortex: The anterior executive and posterior evaluative regions. *Cereb Cortex* 2(6):435–443.
- Carter CS, Botvinick MM, Cohen JD (1999) The contribution of the anterior cingulate cortex to executive processes in cognition. *Rev Neurosci* 10(1):49–57.
- Wager TD, Smith EE (2003) Neuroimaging studies of working memory: A meta-analysis. *Cogn Affect Behav Neurosci* 3(4):255–274.
- Garavan H, Ross TJ, Murphy K, Roche RA, Stein EA (2002) Dissociable executive functions in the dynamic control of behavior: Inhibition, error detection, and correction. *Neuroimage* 17(4):1820–1829.
- Botvinick MM, Cohen JD, Carter CS (2004) Conflict monitoring and anterior cingulate cortex: An update. *Trends Cogn Sci* 8(12):539–546.
- Crick F (1995) *The Astonishing Hypothesis: The Scientific Search for the Soul* (Simon & Schuster, New York).
- Shackman AJ, et al. (2011) The integration of negative affect, pain and cognitive control in the cingulate cortex. *Nat Rev Neurosci* 12(3):154–167.
- Etkin A, Egner T, Kalisch R (2011) Emotional processing in anterior cingulate and medial prefrontal cortex. *Trends Cogn Sci* 15(2):85–93.
- Menon V, Uddin LQ (2010) Saliency, switching, attention and control: A network model of insula function. *Brain Struct Funct* 214(5-6):655–667.
- Legrain V, Iannetti GD, Plaghki L, Mouraux A (2011) The pain matrix reloaded: A saliency detection system for the body. *Prog Neurobiol* 93(1):111–124.
- Rainville P, Duncan GH, Price DD, Carrier B, Bushnell MC (1997) Pain affect encoded in human anterior cingulate but not somatosensory cortex. *Science* 277(5328):968–971.
- Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD (2011) Large-scale automated synthesis of human functional neuroimaging data. *Nat Methods* 8(8):665–670.
- Poldrack RA (2006) Can cognitive processes be inferred from neuroimaging data? *Trends Cogn Sci* 10(2):59–63.
- Wager TD, et al. (2013) An fMRI-based neurologic signature of physical pain. *N Engl J Med* 368(15):1388–1397.
- Chang LJ, Yarkoni T, Khaw MW, Sanfey AG (2013) Decoding the role of the insula in human cognition: Functional parcellation and large-scale reverse inference. *Cereb Cortex* 23(3):739–749.
- Andrews-Hanna JR, Saxe R, Yarkoni T (2014) Contributions of episodic retrieval and mentalizing to autobiographical thought: evidence from functional neuroimaging, resting-state connectivity, and fMRI meta-analyses. *Neuroimage* 91:324–335.
- Poldrack RA, et al. (2012) Discovering relations between mind, brain, and mental disorders using topic mapping. *PLoS Comput Biol* 8(10):e1002707.
- Wendelken C (2014) Meta-analysis: How does posterior parietal cortex contribute to reasoning? *Front Hum Neurosci* 8:1042.
- Ripollés P, et al. (2014) The role of reward in word learning and its implications for language acquisition. *Curr Biol* 24(21):2606–2611.
- Neurosynth Google Group Forum. Available at <https://groups.google.com/forum/#!topic/neurosynth/YLZrwJkqirw>. Accessed May 22, 2015.
- Rosenthal R (1991) *Meta-Analytic Procedures for Social Research* (Sage, Newbury Park, CA), Vol 6.
- Eisenberger NI, Lieberman MD (2004) Why rejection hurts: A common neural alarm system for physical and social pain. *Trends Cogn Sci* 8(7):294–300.
- Bush G, Luu P, Posner MI (2000) Cognitive and emotional influences in anterior cingulate cortex. *Trends Cogn Sci* 4(6):215–222.
- Spunt RP, Lieberman MD, Cohen JR, Eisenberger NI (2012) The phenomenology of error processing: The dorsal ACC response to stop-signal errors tracks reports of negative affect. *J Cogn Neurosci* 24(8):1753–1765.
- Hajcak G, Foti D (2008) Errors are aversive: Defensive motivation and the error-related negativity. *Psychol Sci* 19(2):103–108.
- Inzlicht M, Bartholow BD, Hirsh JB (2015) Emotional foundations of cognitive control. *Trends Cogn Sci* 19(3):126–132.
- Vorobiev V, Govoni P, Rizzolatti G, Matelli M, Luppino G (1998) Parcellation of human mesial area 6: Cytoarchitectonic evidence for three separate areas. *Eur J Neurosci* 10(6):2199–2203.
- Vogt BA, Pandya DN, Rosene DL (1987) Cingulate cortex of the rhesus monkey. I. Cytoarchitecture and thalamic afferents. *J Comp Neurol* 262(2):256–270.
- Craig AD (2003) A new view of pain as a homeostatic emotion. *Trends Neurosci* 26(6):303–307.
- Denton DA, McKinley MJ, Farrell M, Egan GF (2009) The role of primordial emotions in the evolutionary origin of consciousness. *Conscious Cogn* 18(2):500–514.
- Baumeister RF, Leary MR (1995) The need to belong: Desire for interpersonal attachments as a fundamental human motivation. *Psychol Bull* 117(3):497–529.
- Holt-Lunstad J, Smith TB, Layton JB (2010) Social relationships and mortality risk: A meta-analytic review. *PLoS Med* 7(7):e1000316.
- Rotge JY, et al. (2015) A meta-analysis of the anterior cingulate contribution to social pain. *Soc Cogn Affect Neurosci* 10(1):19–27.
- Tataranni PA, et al. (1999) Neuroanatomical correlates of hunger and satiation in humans using positron emission tomography. *Proc Natl Acad Sci USA* 96(8):4569–4574.
- Liotti M, et al. (2001) Brain responses associated with consciousness of breathlessness (air hunger). *Proc Natl Acad Sci USA* 98(4):2035–2040.
- von Leupoldt A, et al. (2009) Dyspnea and pain share emotion-related brain network. *Neuroimage* 48(1):200–206.
- Evans KC, et al. (2002) BOLD fMRI identifies limbic, paralimbic, and cerebellar activation during air hunger. *J Neurophysiol* 88(3):1500–1511.
- Denton D, et al. (1999) Correlation of regional cerebral blood flow and change of plasma sodium concentration during genesis and satiation of thirst. *Proc Natl Acad Sci USA* 96(5):2532–2537.
- Denton D, et al. (1999) Neuroimaging of genesis and satiation of thirst and an interoceptor-driven theory of origins of primary consciousness. *Proc Natl Acad Sci USA* 96(9):5304–5309.
- Farrell MJ, et al. (2008) Effect of aging on regional cerebral blood flow responses associated with osmotic thirst and its satiation by water drinking: A PET study. *Proc Natl Acad Sci USA* 105(1):382–387.
- Yücel M, et al. (2001) Hemispheric and gender-related differences in the gross morphology of the anterior cingulate/paracingulate cortex in normal volunteers: an MRI morphometric study. *Cereb Cortex* 11(1):17–25.
- Foltz EL, White LE, Jr (1962) Pain “relief” by frontal cingulotomy. *J Neurosurg* 19:89–100.
- Johansen JP, Fields HL, Manning BH (2001) The affective component of pain in rodents: Direct evidence for a contribution of the anterior cingulate cortex. *Proc Natl Acad Sci USA* 98(14):8077–8082.
- Barthas F, et al. (2015) The anterior cingulate cortex is a critical hub for pain-induced depression. *Biol Psychiatry* 77(3):236–245.
- Hsu MM, Shyu BC (1997) Electrophysiological study of the connection between medial thalamus and anterior cingulate cortex in the rat. *Neuroreport* 8(12):2701–2707.
- Sikes RW, Vogt BA (1992) Nociceptive neurons in area 24 of rabbit cingulate cortex. *J Neurophysiol* 68(5):1720–1732.
- Yamamura H, et al. (1996) Morphological and electrophysiological properties of ACCx nociceptive neurons in rats. *Brain Res* 735(1):83–92.
- Paus T, et al. (1996) Human cingulate and paracingulate sulci: Pattern, variability, asymmetry, and probabilistic map. *Cereb Cortex* 6(2):207–214.
- Amiez C, et al. (2013) The location of feedback-related activity in the midcingulate cortex is predicted by local morphology. *J Neurosci* 33(5):2217–2228.
- Amiez C, Petrides M (2014) Neuroimaging evidence of the anatomo-functional organization of the human cingulate motor areas. *Cereb Cortex* 24(3):563–578.
- Fornito A, et al. (2008) Variability of the paracingulate sulcus and morphometry of the medial frontal cortex: Associations with cortical thickness, surface area, volume, and sulcal depth. *Hum Brain Mapp* 29(2):222–236.
- Le Provost JB, et al. (2003) Paracingulate sulcus morphology in men with early-onset schizophrenia. *Br J Psychiatry* 182:228–232.

Supporting Information

Lieberman and Eisenberger 10.1073/pnas.1515083112

Affective vs. Sensory Aspects of Pain in dACC

If “pain” is the term most closely associated with dACC activity, what does this finding suggest about the dACC’s function? Is it a dedicated physical pain module? One prominent account of dACC suggests that it is involved in the affective, but not the sensory, aspect of pain processing (12). This account proposes that the dACC supports the distressing part of pain but is not involved in tracking the location or intensity of the pain inputs. Lesion data from humans and animals support this claim (43–45). Moreover, animal research has shown that individual neurons within the dACC have almost no stimulus localization information (46–48). Neurosynth results support this dissociation as well. In addition to term-based reverse inference maps, Neurosynth has broader topic-based maps that use a collection of related terms and weights them by their centrality to the topic.

There are two pain-related topics in the database. For one of these topics (topic no. 80), the top five terms are a mix of generic and more affective pain words (“pain,” “painful,” “stimulation,” “chronic,” “noxious”). For the other (topic no. 131), the top five terms are more focused on the sensory aspects (“somatosensory,” “stimulation,” “tactile,” “touch,” “primary”). Many terms are common to both topics (“pain,” “painful,” “stimulation,” “somatosensory,” “nociceptive,” “sensory,” “perception,” “primary,” “sensation”), but the topics are weighted differently. It is thus of note that the reverse inference map for the somatosensory-focused pain topic does not include activity in the dACC or anterior insula (AI) but instead shows effects in the somatosensory cortex and posterior insula. In contrast, the more affectively focused pain topic includes both the dACC and AI.

Paracingulate Sulcus

There is substantial sulcal variability within the dACC. Although the dACC was historically assumed to consist of the cingulate gyrus and the cingulate sulcus, which sits above it, Paus et al. (49) reported that a second sulcus, the paracingulate sulcus (PCS), is present in a subset of the population and thus extends the dACC further in the dorsal direction. This possible additional sulcus is relevant because, for some individuals, the ventral portion of the SMA/pre-SMA (Fig. 1) may actually be the PCS. The critical question, then, is whether effects we have designated as outside the dACC (e.g., the maximal point of forward inference for the term “dACC”; coordinates 0, 18, 49) might be in the dACC after all.

There is no way to definitively rule out this possibility in the current study. Neurosynth doesn’t have coding for individual participant morphology. Moreover, almost no fMRI studies account for these individual differences. The vast majority of fMRI studies overlook most individual differences in neuroanatomy and depend on the probabilistic neuroanatomy averaged across a group of participants and then on standard atlases that typically don’t take these individual differences into account. Even if we knew definitively that observed effects for terms like “dACC,” “executive,” and “conflict” were from PCS, the current data would still constitute strong data that the majority of the dACC (i.e., the cingulate gyrus and cingulate sulcus) is selective for pain over the various other accounts of dACC function.

However, we think it is unlikely that the Neurosynth results observed in the ventral SMA/pre-SMA were really PCS effects. The statistic typically used to report PCS prevalence is the percentage of individuals who have a PCS of any kind in at least one hemisphere. Across six MRI studies, ~72% of participants met this criterion (42, 49–53). However, three factors reduce the

likelihood that effects observed in the relevant studies in the Neurosynth database are PCS, rather than SMA/pre-SMA.

First, functional activations in this region from individuals with unilateral PCS are likely only resulting from actual PCS 50% of the time (and SMA/pre-SMA the other 50% of the time). Second, there are two structural forms of PCS. The “prominent” form extends through the entire dACC region; however the “present” form begins in the rostral ACC and ends near the anterior border of the dACC. Thus, only the prominent variant of the PCS covers the region in the ventral SMA/pre-SMA under consideration here. Finally, men are significantly more likely than women to have unilateral or bilateral PCS. This gender difference is of consequence because, across PCS morphology studies, the samples are biased toward more males (60%) whereas, in the Neurosynth studies relevant here (e.g., those using the term “dACC”), the samples averaged only 48% male. Thus, population estimates from the morphology studies overestimate the prevalence of the PCS in our Neurosynth sample.

To better estimate the true likelihood that effects in the region we have labeled ventral SMA/pre-SMA are actually from PCS activations, we used data, from a large study ($n = 171$) by Yücel et al. (42), that provide all of the relevant cross-tabulations on how many subjects have prominent or present PCS unilaterally or bilaterally or are missing it altogether. Starting a bit higher than PCS studies in general, 89% of individuals in this study have at least one unilateral PCS of some kind. However, only 60% of participants have prominent PCS, the only form that could produce the activations in question. Moreover, only 16% exhibit bilateral prominent PCS, with 32% showing left unilateral prominent PCS and 12% showing right unilateral prominent PCS. Given that unilateral prominent PCS contribute only a 50% probability of producing observed midline effects, these three variants (bilateral, left unilateral, and right unilateral) suggest a 38% likelihood of observed ventral SMA/pre-SMA effects actually coming from the PCS [i.e., $16\% + 0.5 \times (32\% + 12\%)$].

Finally, men were overrepresented in this sample (58%) and were significantly more likely to show evidence of at least unilateral prominent PCS than women (68% to 50%). After computing the reduced contributions from unilateral prominent PCS, men and women showed 43.5% and 31% likelihoods, respectively, of producing ventral SMA/pre-SMA effects from the PCS. Adjusting for the gender differences across research populations suggests that, in our Neurosynth sample, there is only a 37% chance [i.e., $(\text{male: } 43.5\% \times 48\%) + (\text{female: } 31\% \times 52\%)$] that these effects resulted from PCS tissue and a 63% chance that effects in the region in question came from the SMA/pre-SMA.

Additionally, these six morphology studies (42, 49–53), including the one by Yücel et al. (42), have indicated the existence of a PCS that is left-lateralized. Across these studies, about 35% of participants showed evidence of a prominent PCS in only the left hemisphere whereas only 17% of participants showed evidence of a prominent PCS in only the right hemisphere. If effects we have labeled as SMA/pre-SMA were really PCS, one might expect them to be left-lateralized. Instead, effects tend to either be cleanly bilateral or somewhat right-lateralized.

Across the thousands of participants in the studies examined from the Neurosynth database, a nontrivial number undoubtedly have PCS in the location we have labeled SMA/pre-SMA. Nevertheless, for the reasons given in this section, we think that the effects we observed in this region can more confidently be attributed to the SMA/pre-SMA than to the PCS.

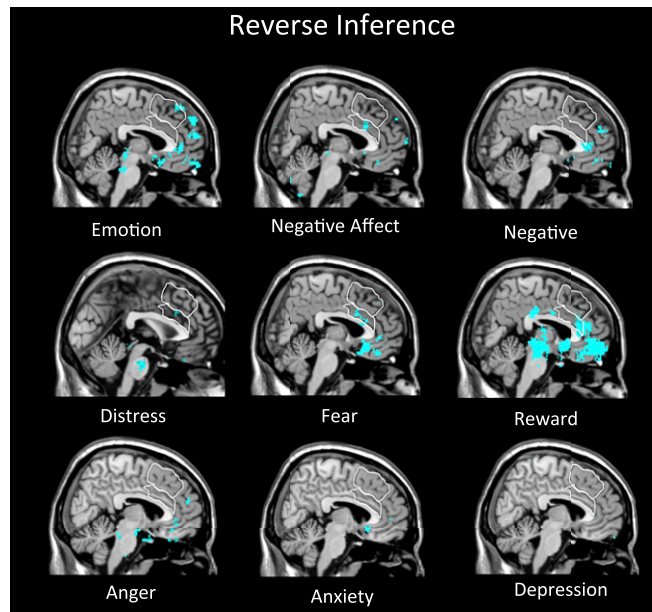


Fig. S1. Neurosynth reverse inference maps for emotion-related terms. Distress-related terms, including "negative affect," "distress," and "fear" yield reliable reverse inference effects in the dACC (but not terms such as "emotion," "negative," "anger," "anxiety," or "depression"). Maps use Neurosynth's standard FDR criterion of $P < 0.01$.

