A Neurophilosophical Investigation of the Limits of fMRI

Kyoko Kamishima[1], Greg Siegle[2]
[1] University of Pittsburgh
[2] University of Pittsburgh School of Medicine Department of Psychiatry

Abstract

This study examines whether neuroimaging can give information about an individual's subjective experiences with emotional stimuli. The first study is a reanalysis of a data set investigating differences in brain activity between stimuli that are personally relevant and not personally relevant to the participants. The second study investigates whether similar relationships between words for individual subjects translate to similar brain activity.

Introduction

New technological advances in neuroscience research have allowed the workings of the human brain to become open to empirical exploration. Active brain areas can be monitored with millimeter voxel resolution, transmitters and receptors can be manipulated, and electrical activity can be recorded from the whole brain or individual neurons. On the other hand, cognitive tasks and tests to assess behavior, memory, and attention allow us to study the workings of the human mind. Pairing these two experimental techniques allows us to overlap data about cognitive and brain activity. However, important connections between brain function and experiences of the mind still remain unlinked. An fMRI scan can monitor transient changes in oxygen consumption to indicate that the amygdala is active when we are scared (11,21), and pharmacological studies show us that drug induced activation of dopaminergic systems in the brain induces a state of euphoria (7,25). We can even distinguish electrical activity from the scalp that characterizes distinct mental activities that we are otherwise unaware of. But what do these measurements really tell us about how our minds work? Correlation studies can only reveal what occurs in the brain during an experience. Lesions or inactivation of brain areas only reveal what region or structure is necessary for specific functions. Manipulations can prove general causation yet the precise mechanisms of this causality between brain activity on the cellular level and cognitive functioning appear intangible. Whether we can ever understand the connection between what occurs and *how* it occurs when considering cognitive functions is unclear. Can neuroscience provide satisfying answers as to how our brains become our minds? If so, how will it explain the emergence of the subjective phenomena that we actually experience?

Neurophilosophy

The link between the brain and the mind represents a black box that researchers and philosophers refer to as the mind-body problem (18). Two theories of mind and brain help to demonstrate opposing views that are neither obvious nor proven. Dualists including Plato and Aristotle asserted that the mind is a nonphysical entity separate from the brain and that intelligence could not be qualified or explained in terms of the physical body (24). Materialists contrarily claim that mental states and operations are equivalent to actual brain states and brain operations, thus mental states are indistinct from the physical states (17). Various other theories of mind and brain including functionalism and behaviorism abound (4). Without explaining the link between mind and brain or at least whether there *is* such a link to be explained at all, phenomena that arise from the brain seem to be almost of another world when we look at the organ. At this point a Cartesian dualistic theory of mind may be just as convincing as a materialist view.

For example, philosophers use the term “qualia” to describe the introspective, subjective, and phenomenal aspects of our mental states. How these qualia relate to brain activity is the crux of the mind-body problem. If our minds are merely functions of our brains, then how can we explain these qualia? How do we translate by looking at fMRI or PET, the neural basis of *feelings* of subjective experiences? For example, it is unclear whether we could use fMRI to explain what it *feels* like to taste chocolate or how we would describe the *experience* of seeing red? We all know what chocolate tastes like, and we all know what it’s like to see red. But does unfolding brain activity that correlates to that experience really give us a satisfying answer as to how these phenomena arise? I look at an apple, and tell you everything that happens when I perceive its redness. When you see an apple, you may assume that you are seeing the same thing. But there is an assumption that we make at this point. I call it red because I have learned to label that experience from the past as seeing red, and you call it red for the same reason. But how can we be certain that what you are experiencing as “red” is the same as what I am experiencing as “red”? (12) We can *assume* that if the same thing is happening in both our retinae, cones, and visual cortices, then the experiences also must be the same. However these states are subjective, perceivable only in the mind of the one who experiences it. There is no objective means to being certain of what someone is actually experiencing. Any conclusion at this point is still an assumption; this obstacle is referred to as the subjectivity barrier (15).

Imagine that in the future, we can monitor the activity of every single neuron in a brain simultaneously and continuously over time. Knowing where and when an action potential occurs, when a specific neurotransmitter is released, what specific receptor was activated, and every signaling cascade or ion flux event that occurs at every neuron in the brain; would that effectively explain how mental phenomena arise? When studying the human motor system, we can understand it logically. A signal initiated from the motor cortex travels down the spinal cord and an action potential leads to calcium release in the muscle. The calcium then triggers the physical movement of sarcomeres in the muscle fibers that lead to an observable effect, muscle contraction. The events can be traced sequentially from the motor cortex to muscle movement, and there is no subjectivity barrier since the effects are observable. In contrast, cognitive aspects are intuitively different. Let’s say someone is shown a grotesque picture of a murdered body. We can infer that most people would be disgusted and possibly a bit scared. But we cannot observe nor quantify these specific internal states. Perhaps a psychopath would actually find it amusing or another person will be reminded of a past incident where their friend was attacked. Assorted physiological measures such as heart rates, respiration rates, skin conductance, facial muscle activity and brain activity can be analyzed. But we cannot at this point trace, as in the example of muscle movement, the emergence of a past memory in the friend or the amusement felt by the psychopath. As brain activity is unraveled, we can identify physiological markers that occur before, during, or after a certain cognitive function, but the result is not visible and is therefore subjective. There is no screen in the brain that shows the scientist a picture imagined in a person’s mind or an emotion felt at any time, and there is nothing that we can deduce with certainty about any person’s subjective experience from brain activity.

Implications for Emotion Research

One area in cognitive neuroscience research where this subjectivity barrier becomes significant is in the study of emotions. Since emotional experiences are very personal, current empirical and statistical analyses don't unravel underlying complexities. Researchers utilize various methods to study emotion: emotional words, movies, fear conditioning, and mood induction are only a few examples. But there seems to be something missing when asking someone to feel a certain emotion while they are lying in a scanner or assuming that a subject will be happy if he or she is shown a picture of a puppy. There are also individual differences in emotional reactivity. The friend is saddened by the memories of her friend induced by the picture of a murdered body whereas I simply feel utter disgust, because our emotional experiences are heavily influenced by our autobiographical memories, personalities, and personal experiences.

Though these nuances have not been well-captured in cognitive neuroscience other disciplines have begun to merge scientific theories of the brain with the more personal contents of memory. Deborah Aschheim is an artist whose work illustrates individual differences in semantic networks of emotion from a similar perspective. Her installation, *On Memory,* includes a room full of personal pictures and captions on the walls each based on an autobiographical memory or related thought. These various nodes are interwoven into a huge network representing how her memories are organized and represented in her conscious mind. This artist rendition of the workings of the mind help to show the contrast between a scientific approach to the study of memory and a more personal and revealing look into the intricacies of our individual autobiographies. The work reminds the scientist that every individual's narrative is an integral facet to the functioning of our minds.

Deborah Aschheim's work also resembles the theories of spreading activation in the semantic networks in the brain (5,6). These theories propose that our minds are networks of interconnected nodes that represent separate, yet related concepts. When one particular node is activated, this activation spreads (called spreading activation) along links to other interconnected concepts. As a result of spreading activation, stimuli may induce activation at other nodes, leading to a conscious experience of interrelated thoughts and memories that differ for each individual (similar to the person being reminded of her friend). Her model here shows an artistic representation of her own networks.

Her two works imply important points to consider. As mentioned, a substantial part of emotion includes the subjective experience that differs between individuals according to our perception of it. Thus, when we are exposed to any stimulus that might contain some relevance to us, brain activity is likely to differ in these cases compared to other stimuli that are relatively neutral or irrelevant. If there are differences, it will indicate the importance of variable reactivity to various stimuli, including emotional stimuli.

Another question to consider is the role of our semantic networks and the effects of spreading activation throughout those networks. How are links between words and memories set up throughout our neural networks? Individual differences in the organization of these networks must be an important aspect in altering our emotional experiences, consequently shaping our identities. Can we see these differences in spreading activation and relate this to subjective experiences across individuals? If these differences are detectable, that would introduce a new empirical approach to studying emotions that can include an aspect of individual variation. Perhaps the subjectivity barrier will no longer be a barrier, but something that can be quantified and analyzed.

Examining individual differences in spreading activation throughout cognitive networks with fMRI

The spreading activation theory proposes that memories are set up into organized networks in the brain. These units of each network are interconnected if they are related to the same concept. Once one concept or element is "activated", this activation spreads throughout the network, activating other interconnected units (1). If the conscious aspects of spreading activation, as conveyed through Deborah Aschheim's work can be overlapped with neural activity, this would indicate that empirical measures of brain activity may have the potential to quantify the conscious and subjective state that one may be experiencing. If we know what a person is thinking and what other memories they are reminded of through this spreading activation, and if we can distinguish neural correlates corresponding to these different networks, perhaps one day it will be possible to tell by looking at brain data what kinds of things a person is thinking or remembering, and maybe even how they are feeling in response. The subjectivity barrier then will no longer be an obstacle to cognitive neuroscience research. Both studies in this paper reanalyze fMRI data to take a first step in investigating these possibilities.

Study 1

The first study is a reanalysis of a data set investigating differences in brain activity between stimuli that are personally relevant and not personally relevant to each of the participants. If stimuli that are personally relevant to a person can be shown to induce brain activity that significantly differs from that induced by non-relevant stimuli, we may say that the brain itself, in addition to the conscious mind is aware that the subjective state that it is in is somehow unique to its person. If brain activity shows that it recognizes this self-related personal aspect, it will indicate that fMRI has the potential to be used to study idiosyncratic aspects of cognition that separate one individual from another. This study investigates whether there are significant differences in brain activity in subjects who are shown words that are personally relevant to them and words that are not personally relevant. The two types of stimuli are hypothesized to show distinct differences in activity across individuals.

Methods

fMRI data collected in a previous experiment (Siegle, Thompson, Carter, Steinhauer, & Thase, 2007) was used for analysis. Only the data from 21 healthy subjects with no history of Axis I disorder or recent health problems including alcohol or drug abuse within the past 6 months were included. 34 images of 3.2-mm slices parallel to the anterior–posterior commissure line were acquired (3-T GE scanner, T2*-weighted images showing BOLD contrast; repetition time 1500msec, echo time 25msec,field of view 24cm, flip angle 60°). Each trial lasted 18 seconds, yielding 12-whole brain images per trial. Words were presented in black on a white projection screen set at a .88° visual angle. Subject responses to tasks were recorded via a Psychology Software Tools glove and response/button sequences were counterbalanced across subjects.

For the task participants were shown a fixation cue for 1sec, which was followed by a positive, negative, or neutral word shown for 200msec. The word was followed by a row of Xs for 10.8sec as a mask. Participants responded by pressing a button, answering whether the word was relevant, somewhat relevant, or not relevant to them or their lives, and were instructed to respond as quickly and accurately as possible. Both participant-generated and normed words were used (Siegle et al 2001, 2002, 2003a, submitted). Normed words balanced for arousal, normed affect, word frequency and length were chosen from the ANEW master list (Siegle 1994), including 10 positive, 10 negative, and 10 neutral words each. Participants also generated 10 positive, 10 negative, and 10 neutral words each that were personally relevant to them prior to the experiment. fMRI analyses were conducted with NeuroImaging Software (NIS) and AFNI. Random-effects whole-brain voxelwise analyses of variance (ANOVAs) using subject as a random factor and personal relevance and scan as fixed factors identified regions with significant personal relevance x scan condition interactions subject to 32 voxel empirically derived contiguity threshold ($p < .005$). Times-series data for each identified ROI were obtained.

Results

Analysis of the fMRI data showed several brain regions with significant personal relevance x scan interactions. ROIs obtained in the analysis located outside grey matter were ignored. Out of the remaining ROIs, two regions showed greater activity in response to personally relevant stimuli compared to non-personally relevant stimuli. Figure 1 shows the anterior cingulate cortex (ACC), the posterior cingulate cortex (PCC) and the retrosplenial cortex (RSC), all showing greater activation in the personally relevant condition.

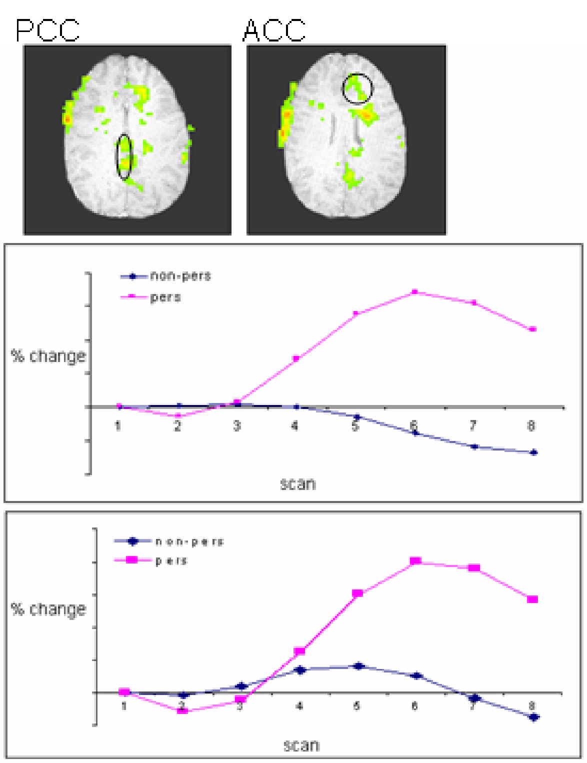

Figure 1

Figure 2

Time course of activation across scans showing differences between conditions. Greater activation in the dorsal ACC was seen in the personally relevant condition.

Greater activation in the left PCC and RSC (BA 23, 30, 31) was also seen in the personally relevant condition.

Discussion

The results of this study are consistent with previous literature attributing cortical midline structures as an integral component in the processing of self-related stimuli. Self-referential or self-related stimuli have been defined as those that are experienced as strongly related to one's own person (14). It has been previously noted that continuous processing of self-related information involves the cortical midline structures (CMS). A meta-analysis by Northoff et al of various studies on self-referential processing demonstrated that cortical midline structures, including the VMPFC, DMPFC, PCC/RSC, and ACC show activity in verbal self-referential processing tasks irrespective of the sensory modalities.

It has been hypothesized that self-referential stimuli are monitored in the supragenual anterior cingulate cortex (SACC including BA 24, 32). The SAC has been implicated in monitoring performance, controlling response selection, and error monitoring (2,3,8,20), and previous studies have concluded that this monitoring function of the SAC shows greater activity in response to self-related stimuli across all sensory modalities. Evaluation of self-referential stimuli has been attributed to the dorsomedial prefrontal cortex (DMPFC including BA9 and BA10). The dorsal CMS, including the DMPFC and SACC have dense connections with the lateral prefrontal cortex, and have been suggested to be involved in the reappraisal and evaluation of self-referential stimuli. The involvement of lateral prefrontal regions has been seen to be a part of reciprocal modulation between self-referential processing and higher order processing of cognitive-emotion interactions (17). The two functions of the SAC and DMPFC are thought to then be integrated in the posterior cingulate cortex (PCC including BA23, 31) where the information may somehow be integrated with the emotional and autobiographical self.

Several imaging studies have also supported a role of the PCC and precuneus in the integration of self-related stimuli with the core sense of self. For example, studies where subjects were asked whether a written word or statement described them or not, or other studies where subjects had to retrieve self-referential information showed activity in the PCC and precuneus. (9,10,13,14). The PCC and RSC have dense connections to the hippocampus, suggesting that these areas may be involved in integrating self-related stimuli within a temporal context, linking them to the domain of autobiographical memory. The RSC has been seen to be involved in the response to emotionally salient information, showing activation to emotionally salient stimuli and in episodic memory tasks. The RSC has dense projections to the parahippocampal cortex and entorhinal cortex and has therefore been hypothesized to mediate interactions between emotion and memory. The PCC and RSC also have reciprocal connections to the DLPFC, possibly suggesting a role in connecting the DLPFC and hippocampal formation.

The experiment supports current theories on how cortical midline structures may function when processing self-relevant stimuli. Differences in brain activity when processing self-relevant information seem to show that we are consistently evaluating stimuli in a way that allows the brain to be "aware" of how anything may be related to ourselves. Since our subjective experiences involving memories are inherently distinct between individuals, the personal relevance of stimuli consequently contributes significantly to the subjectivity barrier. The results show here that the brain activity can distinguish whether something is personally relevant. Therefore in addition to supporting current theories on the processing of these stimuli, these results also show that it is not only our conscious minds that can categorize stimuli in this way, but also our brain activity. This then shows that personal relevance, an inherent part of the subjectivity barrier can be quantified, and shows the potential of fMRI in doing so.

Study 2

The first study showed that differences between personal relevance and non-relevance show similar brain activity consistently across subjects. The consistencies across individuals imply that fMRI can be used to study this one aspect of the subjective experience. However in order to investigate whether fMRI can be used to characterize more personal aspects of subjectivity, individual differences in brain activity between subjects must be considered. In order to utilize the theory of spreading activation in cognitive networks, words were given to participants, and their relationship ratings were used to form graphs representing participants' unique networks. These graphs were therefore used to show how concepts were interrelated within networks for individuals and depict cognitive networks that are realized by the conscious mind of the participants. Words pairs that were associated with each other are within the same semantic network according to the theory. Therefore it is hypothesized that these associated word pairs will have similar patterns of spreading activation, showing a direct correlation between word pair strength and overlapping activation. If patterns of spreading activation resemble each individual's conscious networks, this will indicate that fMRI can be used to study an even more subjective aspect of idiosyncratic, personal experience.

Methods

Four subjects who completed the same task used in Study 1 were given a 40 x 40 word grid with the 30 normed and 10 negative idiosyncratic words that were presented to them during the task in the scanner. They were asked to rate every word with every other word using the grid, rating how closely related the two words were to each other based on personal histories and memories. (The instructions read; "Please rate each pair of words according to how closely the two words are related to each other FOR YOU based on your past experiences of memories of past events and/or relationships. This is different from rating how closely related the words are in meaning. For example, if a pair of words together remind you of something or someone, that pair should be rated as being more closely related FOR YOU.) The ratings ranged from not related at all to very closely related, on a 1 to 5 scale respectively. The ratings were put into SPSS where the proximity ratings were transformed to distance ratings (i.e. pairs that were rated with a 5 being very closely related were changed to 1, 4 to 2, etc.) The new grid was analyzed using multidimensional scaling (ALSCAL model) which projected the distances into a two dimensional space. For one subject, the X and Y axes seemed to scale relative distances in valence and personal relevance. Figure 3 shows this example of one participant's rating grid transformed after multidimensional scaling in SPSS.

Figure 3

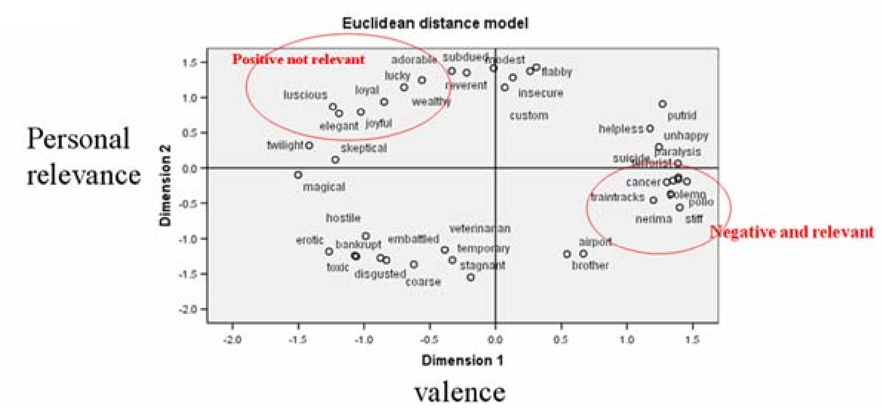

The ratings were then used to show more detailed relationships between all of the words on the multidimensional space. Red lines were used to show strongly related pairs with ratings of 5, and green lines were used to show weaker relationships with ratings smaller than 3. The plots show clusters of words (seen in the multidimensional graphs) and connections within and between words in clusters that represent what would be considered the "interconnected nodes" within a semantic network during spreading activation of brain activity, resembling Deborah Aschheim's model of her own conscious thoughts. Figure 4 shows how different patterns produced from the same words by two individuals show significant variation, depicting the importance of considering idiosyncrasies when investigating emotions.

Figure 4

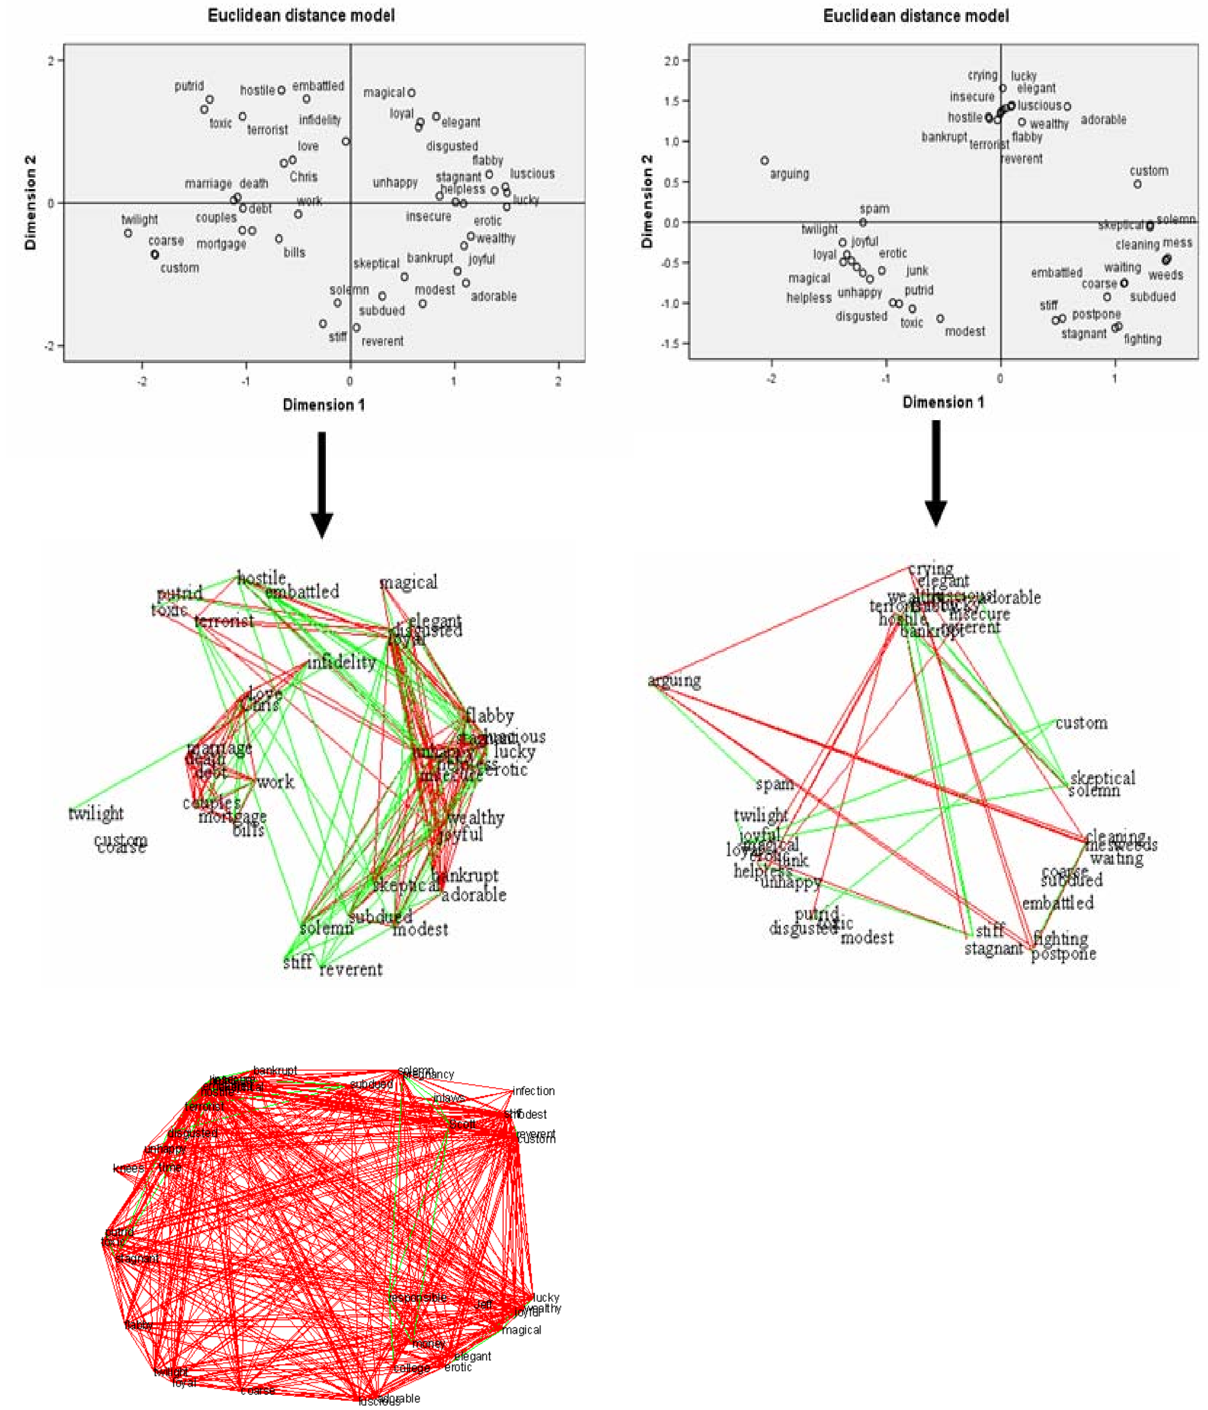

The fMRI data for the three subjects were analyzed. Datasets from the three subjects were transformed to Talairach coordinates using AFNI. Differences in brain activity between each pair of words were calculated for all scans. 18 brain regions including the amygdala, hippocampus, PCC, ACC, transverse, inferior, middle, and superior temporal gyri, and BA areas 9, 24, 29, 30, 31, 32, 33, 45, 46, and 47 were chosen for analysis. These regions were chosen based on previous literature examined in the first study, and also to include temporal areas involved in memory. Each region was first analyzed separately for each subject. Differences in the level of activation across all voxels in the pre-defined regions, between each pair of words were attained. These difference measurements were correlated with strength ratings for each of the word pairs in order to deduce whether strength ratings could predict differences in activation. A second analysis encompassing all of the selected brain regions for each individual was also done to investigate effects of widespread activation throughout relevant regions of the brain.

Results

Analysis of separate brain regions did not show any significant results. Figure 5 shows two examples of results obtained from two selected brain regions, the PCC and BA30. The graphs show the correlation between word pair strengths and the sums of the differences in BOLD response to those word pairs. The multiple regression analysis including all 18 brain regions yielded the last 3 graphs. The graphs in Figure 6 show the relationship between word pair strengths and the predictability of brain activity corresponding to each strength rating. The results show that word pair strengths were able to predict about 10% of the original variance in brain activity.

Figure 5

Graphs show examples of data obtained from analysis of individual regions. Strength ratings between word pairs are represented on the x- axis and the absolute sum of differences in % change of BOLD measurements are along the y-axis.

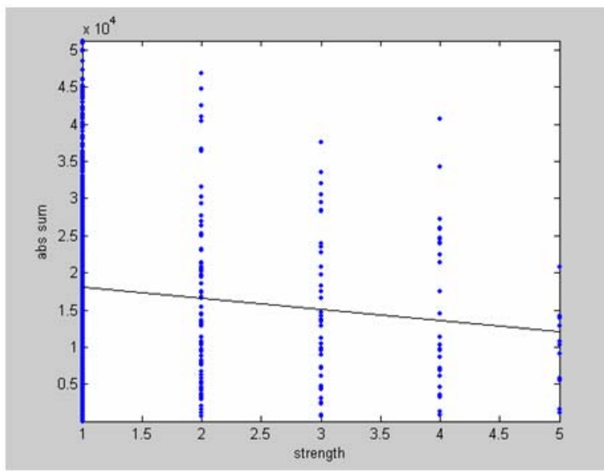

PCC

Rsq = 0.010, F(1,779)=7.771, p=0.005, r=-0.099

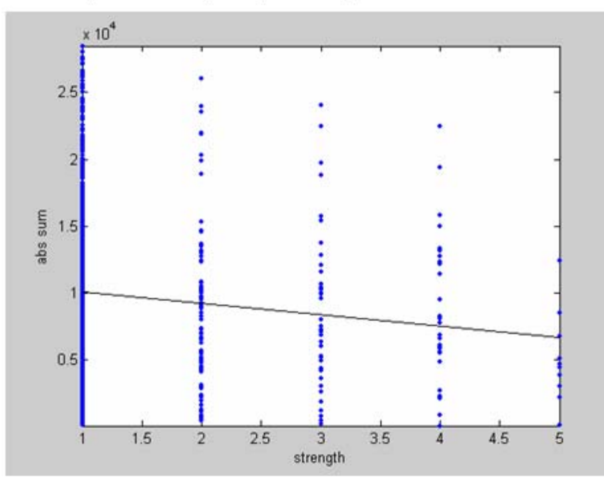

BA30

Rsq = 0.010, F(1,779)=8.169, p=0.004, r=-0.102

Figure 6

Graphs indicating predictability of strength based on brain activity analyzed for all 18 regions. Each graph represents each subject.

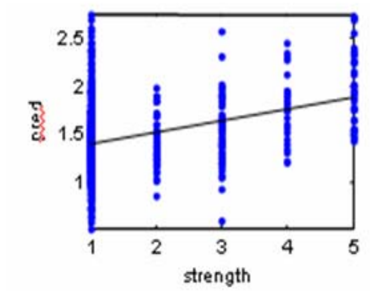

All data: Rsq = 0.119, F(1,779)=104.758, p=0.000

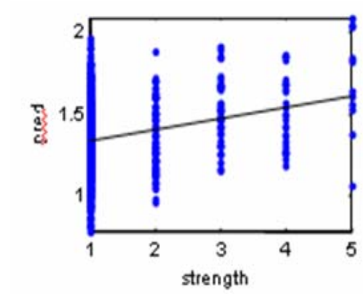

All data: Rsq = 0.068, F(1,779)=56.798, p=0.000

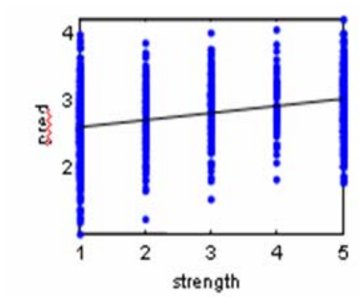

All data: Rsq = 0.106, F(1,779)=92.277, p=0.000

Discussion

Although the first analysis of separate brain regions did not show any significant correlations, the analyses of individual subjects for the 18 regions combined showed slightly stronger results. R squared values show that only about 10% of variance could be predicted, however this effect was seen in all three subjects' data. Despite the weak correlations, evidence of the relationship proves promising for this type of individual analysis.

Both studies reveal that fMRI may have the potential to address some neurophilosophical questions raised earlier in the paper, though analyses accounting for individual differences in semantic network function would have to be different than methods generally used in fMRI. Current fMRI analysis methods usually compile data on many subjects in order to attain regions of interest common to all of the data sets. In order to address some of the neurophilosophical issues, new ways of analysis that focus on the unique activation patterns differing between individuals should be applied. This paper serves as an example for how fMRI analyses using a similar perspective may be attempted in the future to further test the limits in answering similar questions. Although this paper did not investigate the basic questions concerning the causality of cognitive phenomena or the mind body problem, both studies shed some light onto questions concerning the subjectivity barrier. By considering the effects of individual variation in emotion and memory, the studies introduce a new perspective with which to shape future research. The story of the person who is reminded of her attacked friend by a picture provides a good example. The scenario exemplified an obstacle of the subjectivity barrier where it was argued that any person could have any number of subjective responses to the picture of a murdered body, and the elements of this response could not be known for certain due to its inherent subjectivity. However if this new perspective were to be utilized, it is possible that brain activity could eventually be used to infer the actual contents of an individual's experience, in addition to knowing whether a person finds a specific stimulus as personally relevant to them or not. In theory, perhaps neural activity is capable of "mind-reading" on a personal level. Going even further, suppose that a patient with post traumatic brain disorder is troubled by one specific concept or network of memories. We can imagine that some day, if brain activity of cognitive functions can be manipulated on an individual level, the debilitating component of the traumatic memory in the PTSD patient could be singled out and treated. The idea seems absurd at this point, since we still have much to learn about the basic common functions of the human brain. However a novel approach to stressing individual variation in subjective experiences could open a whole new door to the study of cognitive neuroscience.

References

1. Anderson, John Robert. (1983) A spreading activation theory of memory. *Journal of verbal learning and verbal behavior*. 22 (3), 261–295.

2. Bush, G. et al. (2000) Cognitive and emotional influences in anterior cingulate cortex. Trends Cogn. Sci. 4, 215 – 222

3. Carter, C.S. et al. (1998) Anterior cingulate cortex, error detection, and the online monitoring of performance. Science. 280, 747 – 749

4. Churchland, Paul. Matter and Consciousness. Cambridge, MA: MIT Press, 1988

5. Collins, AM. (1975) A spreading activation theory of cognitive processing. Psychological Review. 82, 407

6. Crestani, F. (1997) Application of spreading activation techniques in information retrieval. AI Review. 11(6), 453-482

7. Dackis, C.A. et al. (1985) Pharmacological approaches to cocaine addiction. J Subst. Abuse Treat. 2(3), 139-145

8. Devinsky, O. et al. (1995) Contributions of anterior cingulate cortex to behaviour. Brain. 118, 279 – 306

9. Fink, G.R. et al. (1996) Cerebral representation of one's own past: neural networks involved in autobiographical memory. J. Neurosci. 16, 4275 – 4282

10. Fossati, P. et al. (2003) In search of the emotional self: an FMRI study using positive and negative emotional words. Am. J. Psychiatry. 160, 1938 – 1945

11. Hariri, A.R. et al (2002) Serotonin Transporter Genetic Variation and the Response of the Human Amygdala. Science. 297, 400-403

12. Jackson, F. 1982. Epiphenomenal qualia. *Philosophical Quarterly* 32:127-136.
13. Johnson, S.C. et al. (2002) Neural correlates of self-reflection. Brain. 125, 1808 – 1814

14. Kircher, T.T. et al. (2000) Towards a functional neuroanatomy of self

processing: effects of faces and words. Cognitive. Brain Research. 10, 133-144

15. MacLennan, B. J. (1995) The investigation of consciousness through phenomenology and neuroscience. *Scale in Conscious Experience: Is the Brain Too Important to be Left to Specialists to Study?*, J. King and K. H. Pribram, Eds. Hillsdale: Lawrence Erlbaum, pp. 25–43

16. McGinn, C. (1989) Can we solve the mind-body problem? Mind. 98, 349-366.

17. Nagel, T. (1974) What is it like to be a bat? Nagel, T. (1974) Philosophical Review (1979) [Mortal Questions] Cambridge University Press Cambridge University Press

18. Northoff, G. et al. (2004) Cortical midline structures and the self. Trends in Cog. Sci. 8, no.3

19. Northoff, G. et al. (2006) Self-referential processing in our brain – A meta-analysis of imaging studies on the self. NeuroImage.

20. Paus, T. (2001) Primate anterior cingulate cortex: where motor control,
drive and cognition interface. Nat. Rev. Neurosci. 2, 417 – 424

21. Phelps, E.A et al. (2001) Activation of the left amygdala to a cognitive representation of fear. Nature Neuroscience. 4, no.4

22. Siegle, G.J. et al. (2006) Increased amygdala and decreased dorsolateral prefrontal BOLD responses in unipolar depression: related and independent features

23. Siegle, G. J., Thompson, W., Carter, C. S., Steinhauer, S. R., & Thase, M. E. (2007). Increased amygdala and decreased dorsolateral prefrontal BOLD responses in unipolar depression: Related and independent features. *Biol Psychiatry, 61*(2), 198-209.

24. Stent, G. S. (1998). Epistemic Dualism of Mind and Body. *Proc. Am. Phil. Soc.* 142.4:578–88

25. Wise, R.A. et al. (1985) Brain mechanisms of drug reward and euphoria. Psychiatr. Med. 3(4), 445-460