



UvA-DARE (Digital Academic Repository)

The neural dynamics of fear memory

Visser, R.M.

Publication date

2016

Document Version

Final published version

[Link to publication](#)

Citation for published version (APA):

Visser, R. M. (2016). *The neural dynamics of fear memory*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 6

First steps in using multi-voxel pattern analysis
to disentangle neural processes underlying
generalization of spider fear

Renée M. Visser

Pia Haver

Robert J. Zwieter

H. Steven Scholte

Merel Kindt

Abstract

A core symptom of anxiety disorders is the tendency to interpret ambiguous information as threatening. Using EEG and BOLD-MRI, several studies have begun to elucidate brain processes involved in fear-related perceptual biases, but thus far mainly found evidence for general hypervigilance in high fearful individuals. Recently, multi-voxel pattern analysis (MVPA) has become popular for decoding cognitive states from distributed patterns of neural activation. Here, we used this technique to assess whether aberrant fear generalization is already present during the initial perception and categorization of a stimulus or emerges during the subsequent interpretation of a stimulus. Individuals with low spider fear (LSF, $n = 20$) and high spider fear (HSF, $n = 18$) underwent functional MRI scanning while viewing series of schematic flowers morphing to spiders. Participants were required to indicate for each picture whether they saw a spider, flower or none of the two. In line with previous studies, individuals with high spider fear were more likely to classify ambiguous morphs as spiders than individuals with low spider fear. To our surprise, support vector machine (SVM) classification in 12 functional ROIs did not reveal a clear bias in the classification of morphs in high fearful individuals. On the contrary: response patterns in visual association areas were more likely to be classified as spiders when individuals were *not* afraid of spiders. Although preliminary, these results tentatively suggest that generalization of spider fear is not a perceptual phenomenon, but emerges at a later stage of information processing. Average activation in sensory areas was heightened in individuals with high fear of spiders, independent of stimulus type. Together, these findings support the idea that univariate analysis and multi-voxel pattern analysis tell complementary stories. The combination of these methods may be valuable for disentangling the parallel and semi-independent processes underlying behavior, yet seems to require special design considerations.

Introduction

The ability to recognize threatening stimuli clearly increases the chances of survival. Given that a known threat can take many forms, it is also adaptive to be cautious with other exemplars of the same semantic category that may predict a similar aversive outcome (Mineka, 1992). Stimulus generalization is the mechanism that enables a fast response to novel potentially threatening stimuli, but it can turn into maladaptive behavior when nonthreatening stimuli or contexts are inappropriately treated as harmful. Maladaptive fear generalization is a characteristic of anxiety disorders and post-traumatic stress disorder (Bishop et al., 2015; Kong et al., 2014; Lissek et al., 2005, 2014; Mineka & Zinbarg, 2006) and may even play a causal role in these disorders (Mathews & MacLeod, 2002; Wilson, MacLeod, Mathews, & Rutherford, 2006). For example, while in spider phobia fear responses to real spiders may be debilitating in itself, fear responses to stimuli that more or less resemble the object of fear (e.g., a piece of dust) may interfere most with daily functioning, as phobic individuals find themselves in a permanent state of hypervigilance, avoiding many 'safe' situations (e.g., not eating tomatoes as their insides resemble the legs of a spider). Clarifying which processes enhance fear generalization will ultimately help to answer the fundamental question of why and how people differ in their disposition to develop maladaptive fear behaviors.

Based on decades of animal conditioning research that focused on the perceptual similarity and discriminability of threatening stimuli, it has been implicitly assumed that overgeneralization of fear is a perceptual deficit (Shepard, 1987). In line with this, examples from research in humans show that fearful individuals judge neutral faces as more negative (Bell et al., 2011; Richards et al., 2002), and that individuals with spider phobia more easily see a spider in pictures morphing from flowers to spiders (Kolassa et al., 2007). Only recently, it became evident that fear generalization depends not solely on the physical properties of threatening stimuli, but also on their conceptual properties (Dunsmoor, Martin, & LaBar, 2012; Dunsmoor, Mitroff, & LaBar, 2009; Dunsmoor, White, & LaBar, 2011; Kindt, 2014; Soeter & Kindt, 2012, 2015b). This raises the question whether overgeneralization observed for perceptual cues (such as when a tomato triggers a fear response) is in fact a perceptual process, or instead emerges at a later stage of processing, when the interpretation of a stimulus is guided (biased) by activation of the dominant (fear) network.

As the observed behavior does not allow us to specify whether overgeneralization of fear already occurs during the initial perception and categorization of a stimulus, or emerges at a later stage, it is necessary to go beyond behavioral observations to study the (neural) processes that drive these behaviors. A number of studies have begun to elucidate brain processes involved in fear-related perceptual biases. These studies mainly found heightened sensory sensitivity to all external

stimuli in high fearful individuals, expressed as enhanced early (100 ms) event-related potentials (ERPs) (Frenkel & Bar-Haim, 2011; Kolassa et al., 2007, 2009; Weymar, Keil, & Hamm, 2013) and heightened responses in visual (association) areas to phobogenic objects, often paralleled by heightened responses in the amygdala (Alpers et al., 2009; Dilger et al., 2003; Straube, Mentzel, & Miltner, 2006). These findings are in line with the commonly observed fear-related attentional bias (Bar-Haim, Lamy, Pergamin, Bakermans-Kranenburg, & Van Ijzendoorn, 2007), suggesting that fear facilitates afferent cortical processing in the human visual cortex when individuals search for potential threat. However, heightened sensory sensitivity does not explain how the classification of a stimulus is biased once it is detected. Studies on generalization of fear in healthy individuals have found that varying degrees of perceptual resemblance to a conditioned stimulus elicit graded responses (generalization curves) in the same neurocircuitry that is involved in the acquisition and expression of conditioned fear (i.e., insula, dorsal anterior cingulate cortex; Dymond et al., 2014) and salience processing in general (e.g., the ventral tegmental area; Cha et al., 2014). While in normal fear these graded responses inversely relate to activation in inhibitory brain systems such as the hippocampus and the ventromedial prefrontal cortex, individuals with generalized anxiety seem specifically impaired in recruiting these systems, broadening the range of stimuli to which they respond with fear (Bishop et al., 2015; Cha et al., 2014; Greenberg, Carlson, Cha, Hajcak, & Mujica-Parodi, 2013).

Even though the aforementioned studies provided useful insights into the brain areas that are hyper- or hypoactive in anxiety disorders, a univariate difference in activation per se does not indicate where in the cortical hierarchy ambiguous stimuli are initially being processed as threatening. In contrast, multi-voxel pattern analysis (MVPA) evaluates the information across groups of voxels, to characterize the distinctive neural representation of a stimulus in a certain brain region (Haxby et al., 2001). By training a classifier on neural patterns related to distinct stimulus classes, one can classify patterns related to novel stimuli, providing a more sensitive way to assess the degree to which different stimuli or cognitive states are alike (Haynes & Rees, 2005; Kamitani & Tong, 2005; Kriegeskorte et al., 2008; Norman et al., 2006), or altered by fear (Dunsmoor et al., 2014; Li et al., 2008; Visser et al., 2015, 2013, 2011).

Here, we combined fMRI with an adapted version of the task used by Kolassa and colleagues (2007), to study overgeneralization of fear in individuals with low and high fear of spiders. Based on previous work, we predicted that individuals with high spider fear (HSF) would be more likely to classify ambiguous morphs as spiders than individuals with low spider fear (LSF). Furthermore, we examined in a data-driven manner whether overgeneralization of fear is associated

with functional anomalies in regions traditionally associated with early perception and object identification (Ungerleider & Haxby, 1994), with fear and saliency (Etkin & Wager, 2007; Hermans et al., 2011; Ipser et al., 2013; Seeley et al., 2007) or with higher cognitive processes (Miller & Cohen, 2001). Using support vector machine classification in all cortical and subcortical areas we assessed at what point in the information-processing stream (i.e., ‘low-level’ visual areas, areas associated with salience processing, or ‘higher’ cortical areas associated with decision making) this bias would become apparent.

Methods

Participants

Participants were recruited by means of advertisements in newspapers, social media and the university website. Selection was based on self-reported spider fear as measured by the Spider Phobic Questionnaire (SPQ; Klorman, Weerts, Hastings, Melamed, & Lang, 1974), with scores above 16 representing high spider fear (HSF) and scores below 6 representing low spider fear (LSF). Of the forty-four participants that were initially included, one participant was excluded because of excessive sleepiness, three participants because they did not comply with task instructions, and two participants because of excessive head motion. The final sample included 18 participants in the HSF condition (all female, 3 left-handed, mean = 24.1, \pm 5.9 s.d. yrs. of age), and 20 participants in the LSF condition (14 female, 2 left-handed, mean = 22.9, \pm 1.8 s.d. yrs. of age). Participants earned €20, - for their participation. All participants gave their written informed consent before participating and had normal or corrected-to-normal vision. None of the participants had knowledge of the Chinese language (see materials). Procedures were executed in compliance with relevant laws and institutional guidelines, and were approved by the University of Amsterdam’s ethics committee (2014-CP-3390).

Apparatus and materials

Stimuli. The experiment consisted of one session of fMRI scanning, during which participants performed a task aimed to assess overgeneralization of spider fear (Figure 1a). This task was a modified version of the task used by Kolassa and colleagues (2007), who generously provided part of the stimulus material. This material consisted of schematic morphs that gradually transformed from a flower into a spider by shifting the outlines of the petals until they turned into spider legs (Figure 1b). Three variations existed of this continuum, with each continuum consisting of seven steps, yielding a total of 21 morphs. The presentation of a morph was alternated with the presentation of

an unambiguous picture (Figure 1b), which was either a spider ($n = 7$), a flower ($n = 7$) or a Chinese character ($n = 7$). We collected these unambiguous pictures from the Web, adjusted their luminance, and separated them from their original background. Both the morphs and unambiguous pictures were presented on a grey background.

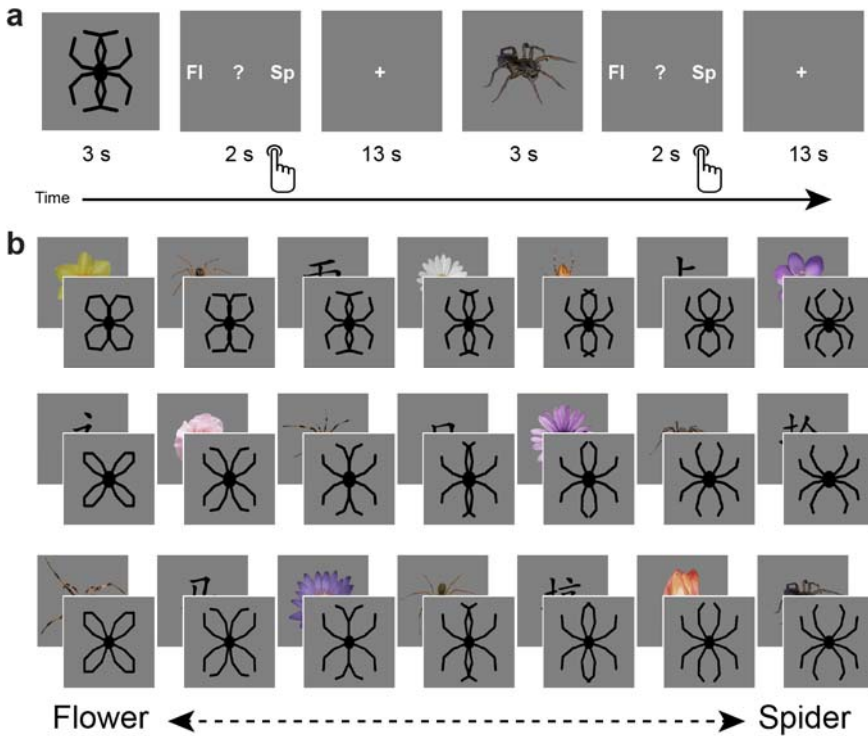


Figure 1 Design. The experiment consisted of one session of fMRI scanning during which a generalization task was performed (a). This task consisted of the presentation of schematic flowers morphing to spiders (generously provided by (Kolassa et al., 2007), intermitted by pictures of spiders, flowers and Chinese characters, to which participants had to respond. Three variations existed of this flower-spider continuum (b). Each variation was presented once, but the fixed order of stimulus presentation was designed in such a way that priming effects could be averaged out: each step of the continuum was once preceded by a flower, once by a spider and once by a Chinese character. Images are not to scale.

Subjective measures. Fear of spiders was assessed with the SPQ (Klorman et al., 1974) and used to select participants. Prior to the experiment, trait anxiety and anxiety sensitivity were assessed with the Trait Anxiety inventory (STAI-T; Spielberger, 1983) and the Anxiety Sensitivity Index (ASI; Peterson & Reiss, 1993) respectively. State anxiety was assessed before and after the scanning procedure with the State Anxiety inventory (STAI-S; Spielberger, 1983).

Image acquisition. Scanning was performed on a 3T Philips Achieva TX MRI scanner using a 32-channel head-coil. Functional data were acquired using a gradient-echo, echo-planar pulse sequence (TR = 2000 ms; TE = 27.63 ms; FA = 76.1°; 37 axial slices with ascending acquisition; 3 × 3 × 3.3 mm voxel size; 80 × 80 matrix; 240 × 133.98 × 240 FoV) and consisted of 415 dynamics. Foam pads minimized head motion, and online motion correction was applied by comparing each recorded volume to the initially recorded volume and adjusting the plane of recording with the displacement. A high-resolution 3D T1-weighted image (TR = 8.30 ms, TE = 3.82 ms, FA = 8°; 1 × 1 × 1 mm voxel size; 240 × 220 × 188 FoV) was additionally collected for anatomical visualization. Stimuli were backward-projected onto a screen that was viewed through a mirror attached to the head-coil.

Pre-processing. fMRI data processing was carried out using FEAT (fMRI Expert Analysis Tool) Version 6.00, part of FSL (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl). Pre-processing included motion correction using MCFLIRT (Jenkinson et al., 2002); slice-timing correction; non-brain removal using BET (Smith, 2002); high-pass temporal filtering ($\sigma = 50$ s), 5 mm spatial filtering and pre-whitening (Woolrich et al., 2001). Registration to high-resolution structural images was carried out using FLIRT (Jenkinson & Smith, 2001) and further refined using FNIRT nonlinear registration (Andersson et al., 2007).

Region of interest selection. Our region of interest (ROI) selection consisted of two steps. First, we conducted a univariate whole-brain analysis to identify clusters that distinguished between unambiguous flower and spider pictures (see section on univariate analysis). The resulting parametric map was then thresholded at $Z > 3.1$ and masked with a whole brain mask created from the Harvard-Oxford cortical and subcortical atlas (part of the FSL software), which excluded brain stem and cerebellum and was thresholded at a probability of $> 10\%$. Next, 12 ROIs were created from clusters consisting of at least 100 adjacent voxels. Using these functional ROIs, we then classified the ambiguous stimuli using a support vector machine (see first two paragraphs of section on multi-voxel pattern analysis).

Second, we conducted a voxel-wise classification analysis in 56 anatomical ROIs (all cortical and subcortical ROIs provided with the Harvard-Oxford cortical and subcortical atlas) to determine in a data-driven manner which voxels best reflected each individual's behavioral responses. All masks were thresholded at $> 25\%$ probability to limit overlap between neighboring regions. Within each ROI, we performed classification analysis to select voxels that best distinguished between unambiguous spiders and flowers (see third paragraph of section on multi-

voxel pattern analysis) and then used this selection for subsequent classification of the ambiguous stimuli.

The reason that we conducted this classification analysis per anatomical ROI, and not on the basis of a whole brain mask, was that we aimed for some regional specificity. Feature selection on the basis of a whole brain mask only revealed clusters in the occipital lobe, ignoring voxels that showed subtler univariate differences between flowers and spiders, but that nevertheless accurately coded for the process of interest.

Experimental design

Upon arrival participants were screened and instructed about the scanning procedure. The experiment started with a structural scan. During functional scanning participants performed the generalization task, viewing morphs (ambiguous) as well as pictures (unambiguous) (Figure 1a, 1b). Each morph was presented once during the task, alternated by the presentation of an unambiguous picture (Figure 1b). Participants were requested to make a response after each stimulus by pressing a button, indicating whether they had seen a spider, a flower, or none of the two (represented by a question mark). With regard to the morphs, we informed participants that the ‘drawings’ they would be seeing would resemble spiders or flowers to a certain degree, while a proportion of these drawings would resemble none of the two. We emphasized that responses to these drawings were purely subjective and that there were no right or wrong answers. Furthermore, we explicitly instructed participants to wait until the stimulus (3 seconds) disappeared and the response screen (2 seconds) was presented. Consequently, reaction times cannot be reliably interpreted, as they do not reflect the initial response to the picture. The response screen (Figure 1a) reminded participants which button to press (left, middle or right), but only the first two letters of the options were shown (‘sp’, ‘fl’, ?), to prevent a fear response to the word ‘spider’ in HSF individuals.

The Chinese characters were included to introduce a clear ‘none-of-the-two’ category, so that responses to the morphs were not biased by response frequencies to the unambiguous stimuli. Stimulus presentation was fixed and was designed in such a way that priming effects could be averaged out: each step of the continuum was once preceded by a flower, once by a spider and once by a Chinese character (Figure 1b). Response buttons and stimulus presentation were counterbalanced across participants. Inter-trial intervals were fixed and relatively long (13 seconds), which seems optimal for single-trial pattern analysis (Visser et al., 2015). The task started with three practice trials (unambiguous pictures of a flower and a spider, and a Chinese character), which were discarded from further analysis. Participants were instructed to pay close attention to the pictures,

even if pictures were unpleasant. Continuous eyetracker-recordings ensured that participants complied with these instructions.

Univariate fMRI analysis

In order to create functional ROIs, and to facilitate interpretation of the results in light of previous fMRI studies on spider fear, we ran a standard voxelwise whole-brain analysis, modeling all trials within a condition as one regressor (10 regressors in total: 7 morph steps (3 per morph), unambiguous flowers (7), unambiguous spiders (7) and Chinese characters (7)) and including 6 motion parameters and temporal derivatives as regressors of no interest. Higher-level mixed-effects analyses were conducted to assess group differences in the contrast of interest, that is, unambiguous spiders > unambiguous flowers. Furthermore, we explored whether there were voxels which' tuning curve followed the gradient of flowers morphing to spiders (i.e., we set up a contrast to test whether there was a linear increase or decrease as function of morphing) and whether these voxels showed overlap with the ones identified using the multivariate approach. Activation was thresholded at $Z > 2.3$ ($Z > 3.1$ for creating the ROIs) and cluster-corrected at $p < 0.05$. Finally, we plotted for each of the functional ROIs the average activation per stimulus type, to examine if the generalization curves mirrored the behavioral data.

Multi-voxel pattern analysis

Single-trial response patterns. Each trial was modeled as a separate regressor in a general linear model (GLM), including six motion parameters as regressors of no interest. The resulting parameter estimates were transformed into t -values to down-weight noisy voxels, by dividing each voxel's parameter estimate by the standard error of that voxel's residual error term after fitting the first-level GLM. In Matlab (version 8.0; MathWorks) we created for each participant, for each ROI a vector containing t -values per voxel for a particular trial. Next, these vectors were used for classification analysis (next two paragraphs).

Classification analysis: functional ROIs. Within each functional ROI, we performed a leave-two-out classification analysis with 1000 iterations, using a one-class support vector machine (SVM) with a linear kernel function (LIBSVM, Chang & Lin, 2011, Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>). With each iteration two unambiguous stimuli (one flower and one spider) were separated as test set, while the other unambiguous stimuli (6 flowers and 6 spiders) were used to train the classifier. Next, the two separated stimuli as well as the 21

ambiguous stimuli were classified, yielding a total number of 23 classifications per iteration (1 flower, 1 spider and 21 morphs). These classifications were averaged over iterations and over the different stimulus types (spider, flower and 7 steps of the flower-spider continuum).

Classification analysis: matching brain and behavior. Aside from the ROI analysis we conducted a voxel-wise classification analysis in 56 anatomical ROIs to determine in a data-driven manner which voxels best reflected each individual's behavioral responses. Hereto, data were again analyzed in a leave-two-out classification analysis with 1000 iterations, except that we now used feature selection. With each iteration two unambiguous stimuli (one flower and one spider) were separated as test set, while the other unambiguous stimuli (6 flowers and 6 spiders) were used to train the classifier and to select the most differentiating voxels. Feature selection thus changed with each iteration: First, z-values were calculated by subtracting the mean t-values over spider trials (6) from the mean t-values over flower trials (6) and dividing the difference by the standard error of the difference. Next, testing was performed using for each individual, within each ROI, the 100 voxels with the highest z-value. Again, this yielded a total number of 23 SVM classifications per iteration (1 flower, 1 spider and 21 morphs). Critically, we now compared the SVM classifications with the responses that the participant made on the ambiguous trials (21 responses, discarding misses and 'none of the two' responses). A percentage was calculated that expressed to what degree SVM classification paralleled behavioral choices. To identify the voxel that accurately paralleled behavior, we assigned a one to the selected features (i.e., the voxels that were involved in that particular classification) if a) the unambiguous spider and flower were correctly classified, and b) if the percentage correctly classified ambiguous stimuli exceeded 90%. Note that 'correct' in this case refers to the correspondence between SVM classification and the behavioral choice. We then divided these 'perfect match'- scores by the total number of iterations. This way we created a whole brain map expressing how often a certain voxel contributed to a near perfect classification ('information map', Chadwick et al., 2012). We then averaged over all individuals, creating a 'perfect match'-map. As the number of voxels that accurately reflected behavior differed substantially across individuals, we created from this map 'perfect-match' ROIs using a minimum cluster size of 100 adjacent voxels and subsequently repeated our SVM analysis (this time without additional feature selection), to visualize the group effects in each ROI. In a way this is double dipping (Kriegeskorte, Simmons, Bellgowan, & Baker, 2009), since the voxels were (across individuals) selected on their match with behavior. However, this ROI analysis was done to explore to what degree the obtained clusters were representative for the groups, that is, truly reflected the behavioral effect.

Statistical analyses

Behavioral data. The behavioral responses, denoted by Y , were dichotomized ($Y = 1$ for "spider", $Y = 0$ for other responses). To account for the nesting of stimulus responses within persons, the data were modeled with a mixed logistic regression model: $P(Y = 1) = \exp(Z)/(1 + \exp(Z))$, where $P(Y = 1)$ denotes the probability that $Y = 1$, and where Z denotes a linear combination of fixed and random effects. Note that these models could also be considered as mixed Rasch models (Kolassa et al., 2007; Rijmen, Tuerlinckx, De Boeck, & Kuppens, 2003). We considered four models: Model 1 consisted of a fixed group effect (LSF vs. HSF, coded as 0 and 1, respectively) and a random person effect. Model 2 consisted of a fixed stimulus effect (the degree to which a stimulus resembles a spider, 7 levels) and a random person effect. Model 3 consisted of a fixed stimulus effect, a fixed group effect, and a random person effect. Model 4 consisted of a fixed stimulus effect, a fixed group effect, a multivariate random person effect with a variance parameter for each group, and a covariance parameter between the two groups.

The model fit was evaluated with AIC and BIC fit statistics, as well as with likelihood ratio tests for nested models. The models were estimated with the statistical software package R, version 3.1.1. (R Core Team, 2014) using the `glmer()` function within the R-package 'lme4' (Bates, Maechler, Bolker, & Walker, 2013; De Boeck et al., 2011).

SVM classification. By iterating training and test sets we obtained normally distributed classification scores for the brain data. Hence, we used parametric tests to assess whether the classification of BOLD-MRI patterns revealed on average more spider classifications in the HSF group, compared to the LSF group. Statistical comparisons of the two groups were performed by between and within-subjects Analysis of Variance (ANOVA), using SPSS (version 21). We specifically tested within each ROI whether there was a main effect of stimulus type (indicating that the region was sensitive to the experimental manipulation), and if significant or trend-significant, whether there was a main effect of group. Predictions were tested while correcting for multiple comparisons (the number of ROIs) by limiting the false discovery rate (FDR; Benjamini & Hochberg, 1995). In case that the assumption of sphericity was violated a Greenhouse-Geisser correction was applied. All p -values are reported two-sided, with the significance level set at $\alpha = 0.05$.

Results

Participant characteristics

Participants in the LSF and HSF group did not differ in trait anxiety and anxiety sensitivity ($p = 0.847$; $p = 0.277$ respectively; Table 1). The (anticipated) confrontation with spider-related material was associated with higher state anxiety in the HSF group, both before ($F_{1, 35} = 6.02$, $p = 0.019$, $\eta_p^2 = 0.15$) and after ($F_{1, 35} = 3.97$, $p = 0.054$, $\eta_p^2 = 0.10$) scanning.

Table 1 Mean values \pm s.d. of self-reported fear of spiders (SPQ), anxiety sensitivity (ASI), state anxiety (STAI-S, pre- and post-scan), and trait anxiety (STAI-T) per group.

	HSF ($n = 18$)	LSF ($n = 20$)
SPQ	21.0 (± 3.0)*	3.2 (± 1.6)*
ASI	8.3 (± 4.2)	10.0 (± 5.0)
STAI-T	35.2 (± 9.5)	35.8 (± 9.0)
STAI-S PRE	36.4 (± 9.1)*	29.3 (± 8.5)*a
STAI-S POST	35.7 (± 13.8)#	28.5 (± 7.4)#a

aBased on 19 participants. SPQ = spider phobia questionnaire; ASI = Anxiety Sensitivity Index. * $p < 0.050$; # $p < 0.080$

Behavioral responses

Figure 2 displays the average proportions in both groups of stimuli identified as spiders, flowers or neither/nor. Table 2 summarizes the fit of the four models as described in the statistical analysis section.

Table 2 Number of estimated parameters (df), log-likelihood (logLik), AIC, and BIC of the four models described in 2.6.1

	df	logLik	AIC	BIC
Model 1	3	-529.72	1065.44	1079.48
Model 2	8	-303.43	622.86	660.32
Model 3	9	-299.53	617.05	659.19
Model 4	11	-298.95	619.89	671.40

The fit statistics indicate that model 3 was the best fitting model. This conclusion was supported by the results of the likelihood ratio tests in Table 3. The fit of model 3 was significantly better than the fit of model 1 and 2, while the fit of model 4 was not significantly better than the fit of model 3.

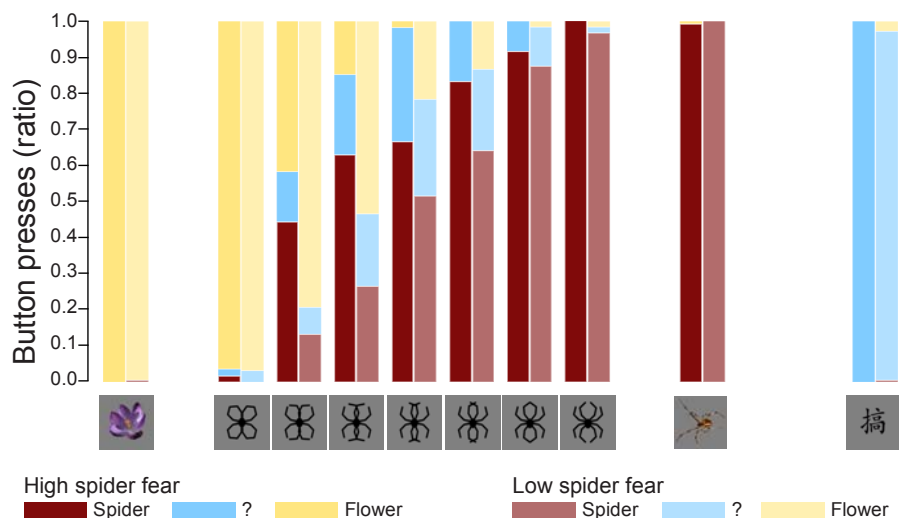


Figure 2 The average proportions in both groups of stimuli identified as spiders, flowers or neither/nor. These results replicate previous findings (Kolassa et al., 2007), suggesting that an individual with high fear of spiders is more likely to classify an ambiguous stimulus as a spider.

Table 3 Results likelihood ratio tests

	-2log(Lik-ratio)	df	p
Model 1 vs. model 3	460.38	6	< 0.0001
Model 2 vs. model 3	7.80	1	0.005
Model 3 vs. model 4	1.16	2	0.56

The parameter estimates of model 3 are displayed in Table 4. The interpretation of the parameters is as follows: since LSF is arbitrarily chosen as reference group, the probability that a randomly selected person with LSF will classify morph 3 as "spider" is $\exp(-0.9872)/(1+\exp(-0.9872)) = 0.271$, while the probability that a randomly selected person with HSF will classify morph 3 as "spider" is $\exp(-0.9872+1.4036)/(1+\exp(-0.9872+1.4036)) = 0.603$. Note that the estimates of the stimulus parameters increase from morph 1 to 7, which implies that the probability that a stimulus will be classified as spider increases from morph 1 to 7. The estimated group effect is statistically significant, $z = 2.889$, $p = 0.004$. These results replicate previous findings (Kolassa et al., 2007), suggesting that an individual with high fear of spiders is more likely to classify an ambiguous stimulus as a spider.

Table 4 Parameter estimates of model 3

Parameter ^a	Estimate	SE
Fixed effects:		
Morph 1	-6.4112	1.1208
Morph 2	-2.0682	0.4191
Morph 3	-0.9872	0.3895
Morph 4	-0.2625	0.3827
Morph 5	0.6497	0.3896
Morph 6	1.9926	0.4402
Morph 7	3.7894	0.6796
Group	1.4036	0.4859
Random effect:		
Intercept person (s.d.)	1.314	

^aEstimates based on the following parameterization: $\text{glmer}(Y \sim -1 + \text{stimulus} + \text{group} + (1|\text{person}), \text{family} = \text{binomial})$

Univariate fMRI results

Whole brain univariate analyses showed typical salience-network activation in response to spider pictures compared to flower pictures (Table 5 & Figure 3). This effect was strongest in individuals with high spider fear, who showed more activation in the salience network as well as visual (association) areas, than individuals with low spider fear (Table 5). These results are in line with previous findings (Alpers et al., 2009; Dilger et al., 2003; Straube et al., 2006). Univariate analysis on activity related to the ambiguous stimuli revealed a cluster in the occipital cortex that responded more to ‘spiderness’ and a cluster in the dorsal paracingulate cortex that responded more to ‘floweriness’. No group differences were observed (Table 6). Finally, Figure 4 shows the average activation per stimulus type, in 12 functional ROIs. Although individuals with high fear of spiders showed more activation than individuals with low fear of spiders in the occipital cortex and a cluster comprising the left amygdala and insula, these effects only reached trend significance ($p = 0.085$ and $p = 0.082$ respectively) and were not specific for ambiguous stimuli (i.e., no effects of morph in any of the ROIs). They therefore seemed to reflect a type of nonspecific hypervigilance rather than an overgeneralization of fear.

Table 5 Brain areas showing differential activation for the unambiguous pictures ($n = 38$)

Brain region (COG)	MNI coordinates			Volume	
	x	y	z	# voxels	Max. Z
Spider > Flower					
Group mean ($n = 38$)					
Saliency network (i.e., frontoinsula cortex, orbitofrontal cortex, dorsal anterior cingulate extending into posterior cingulate cortex, paracingulate cortex, superior frontal gyrus and juxtapositional lobule cortex, temporoparietal junction, amygdala, thalamus, brainstem, cerebellum) and lateral occipital cortex	1	1	-27	69638	8.49
High spider fear ($n = 18$) > Low spider fear ($n = 20$)					
Lingual gyrus, precuneus, intracalcarine cortex	-3	-62	10	14658	4.91
R superior, middle frontal gyrus, precentral gyrus	23	4	54	1354	4.16
R insula, frontal operculum, orbitofrontal cortex	49	16	-5	1308	4.65
R middle temporal gyrus (temporooccipital part), lateral occipital cortex (inferior division), angular gyrus, parietal operculum cortex, supramarginal gyrus	56	-52	12	1137	4.13
Anterior cingulate cortex	1	18	25	1102	3.95
L insula, frontal operculum, orbitofrontal cortex	-39	12	-8	704	3.89
L parietal operculum cortex, supramarginal gyrus	-56	-39	28	520	3.8
Flower > Spider					
Group mean ($n = 38$)					
R Precentral gyrus, R Postcentral gyrus	9	-31	67	473	3.60
High spider fear ($n = 18$) > Low spider fear ($n = 20$)					
No significant clusters					

Whole brain activation ($Z > 2.3$, cluster-corrected at $p < 0.05$), showing clusters of voxels that discriminate between unambiguous spider and flower pictures, and within this contrast, activation that discriminates between groups. Coordinates are in MNI-space and depict for each significant cluster the Center of Gravity (COG). Labels are derived from the Harvard-Oxford cortical and subcortical atlases. L = left; R = right.

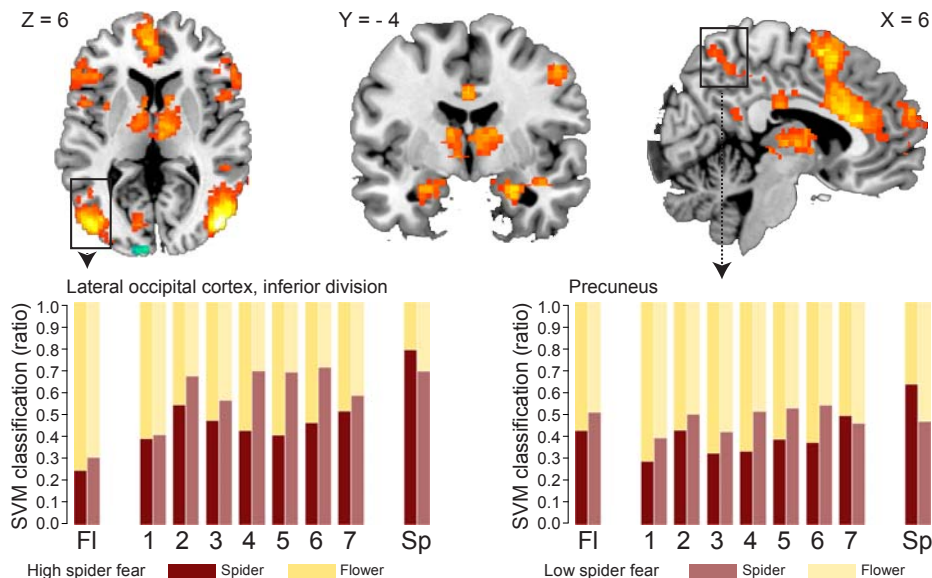


Figure 3 Univariate parametric maps showing clusters of voxels that discriminate between unambiguous flower and spider pictures (top panels, $Z > 3.1$, cluster corrected at $p < 0.05$). Bottom panels show the average SVM classifications in 2 functional ROIs. In the lateral occipital cortex and the precuneus response patterns are more often classified as spider in individuals with low spider fear.

Table 6 Brain areas showing differential activation for the ambiguous pictures ($n = 38$)

Brain region (COG)	MNI coordinates			Volume	
	x	y	z	# voxels	Max. Z
More spider					
Group mean ($n = 38$)					
Occipital cortex	2	-80	8	8591	4.85
High spider fear ($n = 18$) > Low spider fear ($n = 20$)					
No significant clusters					
More flower					
Group mean ($n = 38$)					
Paracingulate gyrus, superior frontal gyrus	4	35	31	1633	3.64
High spider fear ($n = 18$) > Low spider fear ($n = 20$)					
No significant clusters					

Whole brain activation ($Z > 2.3$, cluster-corrected at $p < 0.05$), showing clusters of voxels which' tuning curves followed the gradient of flowers morphing to spiders and vice versa, and within these contrasts, activation that discriminates between groups. Coordinates are in MNI-space and depict for each significant cluster the Center of Gravity (COG). Labels are derived from the Harvard-Oxford cortical and subcortical atlases.

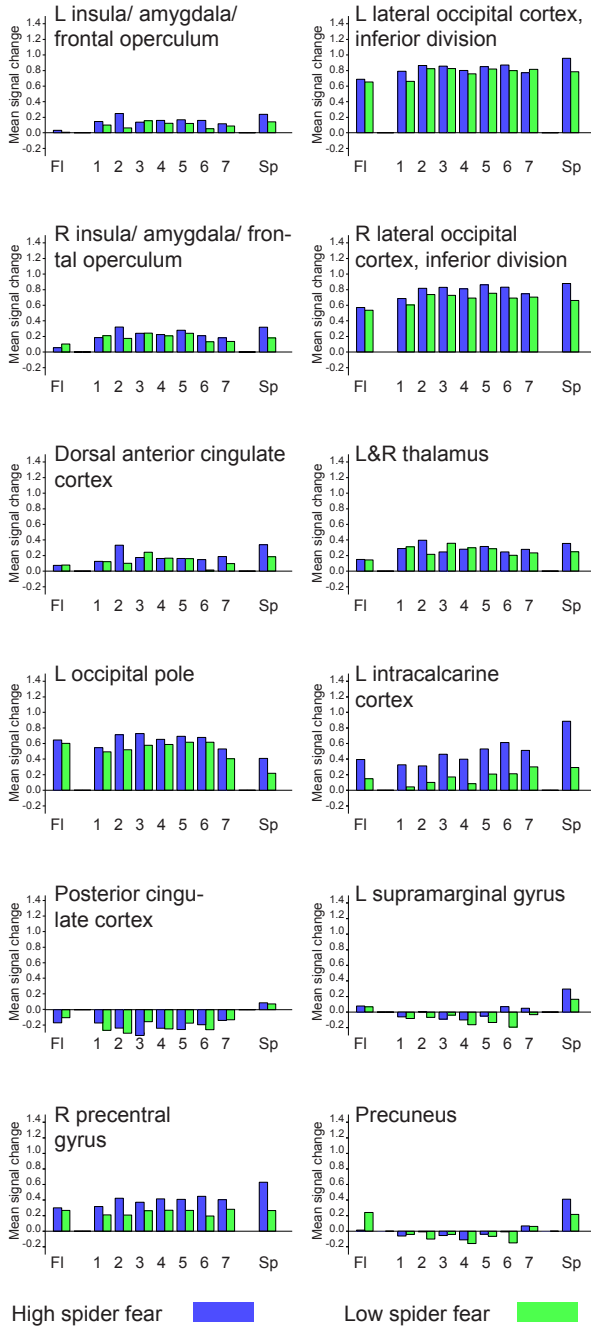


Figure 4 Average activation per stimulus type, in 12 functional ROIs. Although individuals with high fear of spiders showed more activation than individuals with low fear of spiders in the left insula/ amygdala and the intracalcarine cortex, these effects were not specific for ambiguous stimuli, and therefore seems to reflect a type of nonspecific hypervigilance, instead of overgeneralization of fear. L = left; R = right.

Classification results

SVM classification: functional ROIs. SVM classification in 12 functional ROIs did not reveal group differences in the proportion of response patterns classified as spider (Table 7), except for small (uncorrected) effects in the left lateral occipital cortex and the precuneus. However, the effects were in the opposite direction as hypothesized, given that response patterns were more likely to be classified as spiders when individuals were not afraid of spiders (Figure 3). This reversed pattern was visible in all visual association areas, but not in the primary visual cortex. An exploratory analysis showed that classification of *unambiguous* spiders was significantly higher in the HSF group than in the LSF group in a number of regions (ROI 3-5, 11-12, see Table 7, all p s < 0.018), indicating that these pictures elicited more generalized responses when individuals were afraid of spiders (Watts & Dagleish, 1991). It is noteworthy that the classification did not change substantially when we corrected for average activation, which we did by subtracting the signal per trial, averaged across voxels within that ROI (i.e., preserving the spatial pattern, but scaling the activation so that every trial's mean activation in a particular ROI was zero).

Table 7 Summary of statistics of the support vector machine classification ($n = 38$), in 12 functional ROIs

Brain region	# voxels	Main effect of morph (7) (within-subject)			Main effect of group (2) (between-subject)		
		<i>F</i>	<i>p</i>	η_p^2	<i>F</i>	<i>p</i>	η_p^2
ROI 1: L amygdala, anterior insula, inferior temporal gyrus, anterior division, inferior frontal gyrus, frontal operculum, frontal pole	3110	1.86	0.089	0.049	0.296	0.590	0.01
ROI 2: L lateral occipital fusiform cortex, lateral occipital cortex inferior division, occipital fusiform gyrus	4542	3.92	0.001	0.10	5.15 ^a	0.029 ^a	0.13 ^a
ROI 3: R amygdala, anterior insula, inferior temporal gyrus, anterior division, inferior frontal gyrus, frontal operculum, frontal pole	3090	2.42	0.040	0.06	1.97	0.169	0.05
ROI 4: R lateral occipital fusiform cortex, lateral occipital cortex inferior division, occipital fusiform gyrus	4806	5.26	<0.0005	0.13	1.57	0.218	0.04
ROI 5: Dorsal anterior cingulate cortex, paracingulate cortex, superior frontal gyrus, juxtapositional lobule cortex	5044	2.83	0.011	0.07	1.40	0.244	0.04
ROI 6: L & R thalamus	1125	0.96	0.451	0.03	NT	NT	NT
ROI 7: L Occipital pole	113	2.27	0.052	0.06	0.03	0.870	0.00
ROI 8: L intracalcarine cortex	133	6.11	<0.0005	0.15	0.12	0.734	0.00
ROI 9: Posterior cingulate cortex	122	0.84	0.540	0.02	NT	NT	NT
ROI 10: L supramarginal gyrus	397	3.23	0.005	0.08	0.85	0.361	0.02
ROI 11: R precentral gyrus	215	0.37	0.900	0.01	NT	NT	NT
ROI 12: Precuneus	301	2.26	0.039	0.06	3.88 ^a	0.056 ^a	0.10 ^a

All significant values ($p < 0.05$) are in italics, and those that reach FDR-corrected significance are in bold (corrected for 12 ROIs). NT = not tested: Areas without significant main effect of morph are not tested for group effects. ^aGroup effect caused by more classifications as spider in the low spider fear group. L = left; R = right.

SVM classification: matching brain and behavior. The voxel-wise classification analysis revealed four regions that matched individual behavior (Table 8 and Figure 5). However, none of the clusters in this network showed a significant effect of group, indicating that the match between brain and behavior that was found in some individuals was not representative of the entire group. The region that most closely resembled behavior was a region in the postcentral gyrus (Figure 5). In this cluster more spider classifications were made in the HSF group compared to the LSF group.

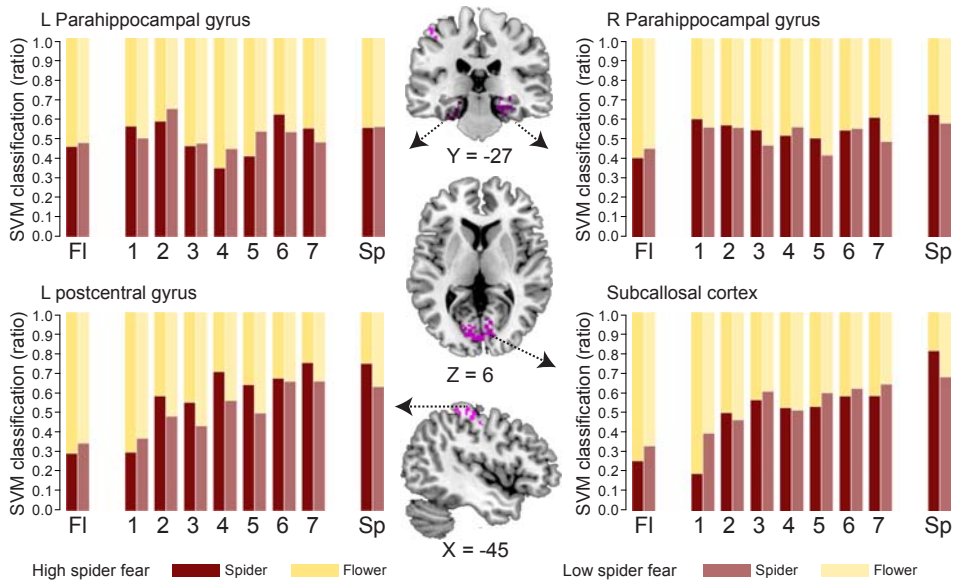


Figure 5 Regions in which support vector machine classification overlapped with behavior in some, but not all individuals. Graphs depict the average SVM classifications per region. L = left; R = right.

Table 8 Summary of statistics of the support vector machine classification ($n = 38$), in 4 ROIs created on match with behavior

Brain region	# voxels	Main effect of morph (7) (within-subject)			Main effect of group (2) (between-subject)		
		<i>F</i>	<i>P</i>	η_p^2	<i>F</i>	<i>P</i>	η_p^2
L parahippocampal gyrus	111	4.11	0.001	0.10	0.08	0.776	0.00
R parahippocampal gyrus, amygdala, hippocampus	191	1.21	0.304	0.03	NT	NT	NT
Subcallosal cortex	440	8.27	<0.0005	0.19	0.60	0.443	0.02
L Postcentral gyrus	123	10.22	<0.0005	0.22	2.17	0.150	0.06

All significant values ($p < 0.05$) are in italics, and those that reach FDR-corrected significance are in bold (corrected for 4 ROIs). NT = not tested: Areas without significant main effect of morph are not tested for group effects. L = left; R = right.

Discussion

The aim of the present study was to examine how fear influences the processing of ambiguous stimuli. Specifically, we used support vector machine (SVM) classification in a data-driven manner to determine at what stage in the information-processing sequence spider fear biases the resolution of ambiguity. In line with previous findings (Kolassa et al., 2007), individuals with high spider fear were more likely to classify ambiguous morphs as spiders than individuals with low spider fear. Unexpectedly, support vector machine (SVM) classification in 12 functional ROIs did not resemble behavior in high fearful individuals. On the contrary: response patterns in visual association areas related to ambiguous stimuli were more likely to be classified as spiders when individuals were *not* afraid of spiders. An exploratory whole brain search - using feature selection and SVM classification - identified a small cluster of voxels in the postcentral gyrus that matched the behavioral results, in some, but not all individuals.

Although preliminary, these results tentatively suggest that overgeneralization of fear is not a perceptual phenomenon, but emerges at a later stage of information processing, which is in line with recent findings in social phobia (Riwkes, Goldstein, & Gilboa-Schechtman, 2015). This notion, however, does not explain the striking dissociation between SVM classification in visual association areas and behavior. Not only do these behavioral and neural levels of responding diverge, they even show opposite effects. This may point at methodological limitations of the current paradigm. Unlike most studies, which used either pictures of phobogenic material (Dilger et al., 2003) or schematic morphs (Kolassa et al., 2007), we alternately presented morphs and pictures of real spiders. In the SVM classification, we trained a classifier on the unambiguous pictures to classify the schematic morphs. This seemed the most ecologically valid approach, as we were interested in the degree to which the morphs would resemble 'true' spiders or flowers, not in the similarity between schematic morphs per se. However, by mixing the two types of stimuli, we may have unintentionally deflated the valence of the spider drawings, skewing the classification of these stimuli. Although these morphs could have been a valid representation of the fearful category if presented alone, the (anticipated) presence of stimuli that are substantially more arousing may have turned the morphs into relatively safe stimuli. For individuals without spider fear both classes of unambiguous stimuli were virtually neutral, so training a classifier on these stimuli showed a linear increase in the likelihood that an ambiguous morph is classified as a spider. For individuals with high spider fear, however, the classifier is trained on two very distinct categories: one neutral and one highly emotional. This is also evident from the fact that these unambiguous categories are better classified - thus more distinct - in high fearful individuals than in low fearful individuals. In this scenario, a relatively neutral morph is

classified as a flower, probably not on the basis of perceptual or even conceptual similarity, but on the absence of a strong emotional response. Between-run classification, with ambiguous and unambiguous stimuli presented in separate runs, could partially solve this problem: without the continuous anticipation of the (terrifying) unambiguous spider pictures, the spider drawings could be semantically categorized as spiders, instead of merely being categorized as 'relatively safe'. Yet, even in separate runs, the unambiguous spider pictures will undoubtedly elicit a much stronger emotional response than the morphs, which may still hamper a balanced classification of stimulus categories if the unambiguous stimuli were to be used for training. These results exemplify that it is methodologically challenging to classify emotional and neutral stimuli in groups that differ in their judgment of emotionality.

Even though the design may have skewed the classification of morphs in many areas of the brain, participants did eventually indicate that they had seen a spider. Given that every behavioral response must have a neural correlate, one would assume that it should be possible with SVM classification to detect an area that codes for this behavior. Although we observed a cluster in the postcentral gyrus that closely matched the behavioral results, this result should be interpreted with caution. First, we did not have prior hypotheses about this area, and its role in decision-making is not immediately obvious. Second, group effects did not reach statistical significance, indicating that there was a substantial amount of inter-individual variability in the degree to which the voxels in this area mirrored behavior. Third, and more importantly, the kind of task may not have been optimal for detecting subtle effects. While overgeneralization of fear was observed at the level of the group, individual behavior consisted of three discrete responses per stimulus type, yielding a proportion of stimuli classified as flower or spider. Thus, individual behavior was not necessarily representative of the group average. Although we based our behavioral task on previous research (Kolassa et al., 2007), future research should explore the possibility of having participants evaluate ambiguous stimuli on a continuous scale, which seems a more sensitive approach for linking individual behavior to brain data. Furthermore, the number and spacing of stimuli used in this experiment were based on previous studies (Visser et al., 2015, 2013, 2011), where pattern analysis was applied to distinguish face and house stimuli (Epstein & Kanwisher, 1998; Haxby et al., 2001) and to assess the effects of Pavlovian conditioning (a strong manipulation). The present manipulation - gradually morphing a flower into a spider - was presumably subtler and therefore required a greater number of trials. Increasing the power and using different behavioral measures would likely yield stronger effects and would open up avenues for model-based searchlights, using (continuous) representational similarity analysis (Kriegeskorte, 2011; Kriegeskorte et al., 2008) instead of the rather coarse

(dichotomous) classification analysis. Approaches like these could further elucidate how initial perceptual evaluation and subsequent conceptual information is eventually combined into deliberate, yet biased, decision-making (Blanchette & Richards, 2010).

Average activation in visual association areas was heightened in individuals with high fear of spiders, independent of stimulus type. As this effect was not specific for ambiguous stimuli, it seemed to reflect a type of nonspecific hypervigilance rather than overgeneralization of fear. Heightened sensory sensitivity fits well with the commonly observed attentional bias in fearful individuals (Bar-Haim et al., 2007). Together, these findings support the idea that univariate analysis and multi-voxel pattern analysis tell complementary stories (Jimura & Poldrack, 2012). The combination of these methods could potentially be valuable for disentangling the parallel and semi-independent processes that build up to behavior, especially when they tap into dissociating components of complex neural systems. However, special attention should be paid to optimizing fMRI designs in such a way that they enable both types of analysis.