

# Shared states: using MVPA to test neural overlap between self-focused emotion imagery and other-focused emotion understanding

Suzanne Oosterwijk,<sup>1,2</sup> Lukas Snoek,<sup>3</sup> Mark Rotteveel,<sup>1,2</sup>  
Lisa Feldman Barrett,<sup>4,5,6</sup> and H. Steven Scholte<sup>2,3</sup>

<sup>1</sup>Department of Social Psychology, University of Amsterdam, The Netherlands, <sup>2</sup>Amsterdam Brain and Cognition Centre, Amsterdam, The Netherlands, <sup>3</sup>Department of Brain and Cognition, University of Amsterdam, The Netherlands, <sup>4</sup>Department of Psychology, Northeastern University, Boston, MA, USA, <sup>5</sup>Athinoula A. Martinos Center for Biomedical Imaging, and <sup>6</sup>Psychiatric Neuroimaging Division, Department of Psychiatry, Massachusetts General Hospital and Harvard Medical School, Charlestown, MA, USA

Correspondence should be addressed to Suzanne Oosterwijk, Department of Social Psychology, University of Amsterdam, The Netherlands. E-mail: s.oosterwijk@uva.nl.

Suzanne Oosterwijk and Lukas Snoek contributed equally to this work.

## Abstract

The present study tested whether the neural patterns that support imagining ‘performing an action’, ‘feeling a bodily sensation’ or ‘being in a situation’ are directly involved in understanding *other people’s* actions, bodily sensations and situations. Subjects imagined the content of short sentences describing emotional actions, interoceptive sensations and situations (self-focused task), and processed scenes and focused on *how* the target person was expressing an emotion, *what* this person was feeling, and *why* this person was feeling an emotion (other-focused task). Using a linear support vector machine classifier on brain-wide multi-voxel patterns, we accurately decoded each individual class in the self-focused task. When generalizing the classifier from the self-focused task to the other-focused task, we also accurately decoded whether subjects focused on the emotional actions, interoceptive sensations and situations of *others*. These results show that the neural patterns that underlie self-imagined experience are involved in understanding the experience of other people. This supports the theoretical assumption that the basic components of emotion experience and understanding share resources in the brain.

**Key words:** multi-voxel pattern analysis; mentalizing; simulation; emotion

## Introduction

To navigate the social world successfully it is crucial to understand other people. But how do people generate meaningful representations of other people’s actions, sensations, thoughts and emotions? The dominant view assumes that representations of other people’s experiences are supported by the same neural systems as those that are involved in generating

experience in the self (e.g. Gallese et al., 2004; see for an overview Singer, 2012). We tested this principle of self-other neural overlap directly, using multi-voxel pattern analysis (MVPA), across three different aspects of experience that are central to emotions: actions, sensations from the body and situational knowledge.

In recent years, evidence has accumulated that suggests a similarity between the neural patterns representing the self and

Received: 4 June 2016; Revised: 23 February 2017; Accepted: 15 March 2017

© The Author (2017). Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

others. For example, a great variety of studies have shown that observing actions and sensations in other people engages similar neural circuits as acting and feeling in the self (see for an overview Bastiaansen et al., 2009). Moreover, an extensive research program on pain has demonstrated an overlap between the experience of physical pain and the observation of pain in other people, utilizing both neuroimaging techniques (e.g. Lamm et al., 2011) and analgesic interventions (e.g. Rütgen et al., 2015; Mischkowski et al., 2016). This process of ‘vicarious experience’ or ‘simulation’ is viewed as an important component of empathy (Carr et al., 2003; Decety, 2011; Keysers and Gazzola, 2014). In addition, it is argued that mentalizing (e.g. understanding the mental states of other people) involves the same brain networks as those involved in self-generated thoughts (Uddin et al., 2007; Waytz and Mitchell, 2011). Specifying this idea further, a constructionist view on emotion proposes that both emotion experience and interpersonal emotion understanding are produced by the same large-scale distributed brain networks that support the processing of sensorimotor, interoceptive and situationally relevant information (Barrett and Satpute, 2013; Oosterwijk and Barrett, 2014). An implication of these views is that the representation of self- and other-focused emotional actions, interoceptive sensations and situations overlap in the brain.

Although there is experimental and theoretical support for the idea of self-other neural overlap, the present study is the first to directly test this process using MVPA across three different aspects of experience (i.e. actions, interoceptive sensations and situational knowledge). Our experimental design consisted of two different tasks aimed at generating self- and other-focused representations with a relatively large weight given to either action information, interoceptive information or situational information.

In the *self-focused* emotion imagery task (SF-task) subjects imagined performing or experiencing actions (e.g. *pushing someone away*), interoceptive sensations (e.g. *increased heart rate*) and situations (e.g. *alone in a park at night*) associated with emotion. Previous research has demonstrated that processing linguistic descriptions of (emotional) actions and feeling states can result in neural patterns of activation associated with, respectively, the representation and generation of actions and internal states (Pulvermüller and Fadiga, 2010; Oosterwijk et al., 2015). Furthermore, imagery-based inductions of emotion have been successfully used in the fMRI scanner before (Wilson-Mendenhall et al., 2011; Oosterwijk et al., 2012), and are seen as robust inducers of emotional experience (Lench et al., 2011). In the *other-focused* emotion understanding task (OF-task), subjects viewed images of people in emotional situations and focused on actions (i.e. *How does this person express his/her emotions?*), interoceptive sensations (i.e. *What does this person feel in his/her body*) or the situation (i.e. *Why does this person feel an emotion?*). This task is based on previous research studying the neural basis of emotion oriented mentalizing (Spunt and Lieberman, 2011).

With MVPA, we examined to what extent the SF- and OF-task evoked similar neural patterns. MVPA allows researchers to assess whether the neural pattern associated with one set of experimental conditions can be used to distinguish between another set of experimental conditions. This relatively novel technique has been successfully applied to the field of social neuroscience in general (e.g. Gilbert et al., 2012; Brosch et al., 2013; Parkinson et al., 2014), and the field of self-other neural overlap in particular. For example, several MVPA studies recently assessed whether experiencing pain and observing pain

in others involved similar neural patterns (Corradi-Dell’Acqua et al., 2016; Krishnan et al., 2016). Although there is an ongoing discussion about the specifics of shared representation in pain based on these MVPA results (see for an overview Zaki et al., 2016), many authors emphasize the importance of this technique in the scientific study of self-other neural overlap (e.g. Corradi-Dell’Acqua et al., 2016; Krishnan et al., 2016).

MVPA is an analysis technique that decodes latent categories from fMRI data in terms of multi-voxel patterns of activity (Norman et al., 2006). This technique is particularly suited for our research question for several reasons. First of all, although univariate techniques can demonstrate that tasks activate the same brain regions, only MVPA can statistically test for shared representation (Lamm and Majdandzić, 2015). We will evaluate whether multivariate brain patterns that distinguish between mental events in the SF-task can be used to distinguish, above chance level, between mental events in the OF-task. Second, MVPA analyses are particularly useful in research that is aimed at examining distributed representations (Singer, 2012). Based on our constructionist framework, we indeed hypothesize that the neural patterns that will represent self- and other focused mental events are distributed across large-scale brain networks. To capture these distributed patterns, we used MVPA in combination with data-driven univariate feature selection on whole-brain voxel patterns, instead of limiting our analysis to specific regions-of-interest (Haynes, 2015). And third, in contrast to univariate analyses that aggregate data across subjects, MVPA can be performed within-subjects and is thus able to incorporate individual variation in the representational content of multivariate brain patterns. In that aspect within-subject MVPA is sensitive to individual differences in how people imagine actions, sensations and situations, and how they understand others. In short, for our purpose to explicitly test the assumption that self and other focused processes share neural resources, MVPA is the designated method.

We tested the following two hypotheses. First, we tested whether we could classify *self-imagined* actions, interoceptive sensations and situations above chance level. Second, we tested whether the multivariate pattern underlying this classification could also be used to classify the how, what and why condition in the *other-focused* task.

## Method

### Subjects

In total, we tested 22 Dutch undergraduate students from the University of Amsterdam (14 females;  $M_{\text{age}} = 21.48$ ,  $s.d._{\text{age}} = 1.75$ ). Of those 22 subjects, 13 subjects were tested twice in 2 sessions about 1 week apart. Half of those sessions were used for the model optimization procedure. The other half of the sessions, combined with an additional nine subjects (who were tested only once), constituted the model validation set (see Model optimization procedure section). In total, two subjects were excluded from the model validation dataset: one subject was excluded because there was not enough time to complete the experimental protocol and another subject was excluded due to excessive movement (>3 mm within data acquisition runs).

All subjects signed informed consent prior to the experiment. The experiment was approved by the University of Amsterdam’s ethical review board. Subjects received 22.50 euro per session. Standard exclusion criteria regarding MRI safety

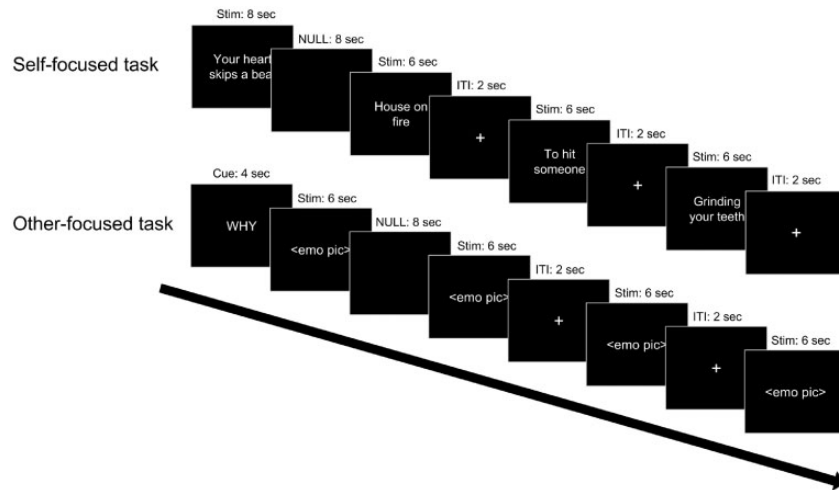


Fig. 1. Overview of the self-focused and other-focused task.

were applied and people who were on psychopharmacological medication were excluded a priori.

### Experimental design

**Self-focused emotion imagery task.** The self-focused emotion imagery task (SF-task) was created to preferentially elicit self-focused processing of action, interoceptive or situational information associated with emotion. Subjects processed short linguistic cues that described actions (e.g. *pushing someone away; making a fist*), interoceptive sensations (e.g. *being out of breath; an increased heart rate*), or situations (e.g. *alone in a park at night; being falsely accused*) and were instructed to imagine performing or experiencing the content. The complete instruction is presented in the Supplementary Materials; all stimuli used in the SF-task are presented in Supplementary Table S1. Linguistic cues were selected from a pilot study performed on an independent sample of subjects ( $n=24$ ). Details about this pilot study are available on request. The descriptions generated in this pilot study were used as qualitative input to create short sentences that described actions, sensations or situations that were associated with negative emotions, without including discrete emotion terms. The cues did not differ in number of words, nor in number of characters ( $F < 1$ ).

The SF-task was performed in two runs subsequent to the other-focused task using the software package Presentation (Version 16.4, [www.neurobs.com](http://www.neurobs.com)). Each run presented 60 sentences on a black background (20 per condition) in a fully randomized event-related fashion, with a different randomization for each subject. Note that implementing a separate randomization for each subject prevents inflated false positive pattern correlations between trials of the same condition, which may occur in single-trial designs with short inter-stimulus intervals (Mumford et al., 2014). A fixed inter-trial-interval of 2 s separated trials; 12 null-trials (i.e. a black screen for 8 s) were mixed with the experimental trials at random positions during each run (see Figure 1).

**Other-focused emotion understanding task.** The other-focused emotion understanding task (OF-task) was created to preferentially elicit other-focused processing of action, interoceptive or situational information associated with emotion. Subjects

viewed images of people in negative situations (e.g. a woman screaming at a man, a man held at gunpoint). A red rectangle highlighted the face of the person that the subjects should focus on to avoid ambiguity in images depicting more than one person. Image blocks were preceded by a cue indicating the strategy subjects should use in perceiving the emotional state of the people in the images (Spunt and Lieberman, 2011). The cue *How* instructed the subjects to identify actions that were informative about the person's emotional state (i.e. *How* does this person express his/her emotions?). The cue *What* instructed subjects to identify interoceptive sensations that the person could experience (i.e. *What* does this person feel in his/her body). The cue *Why* instructed subjects to identify reasons or explanations for the person's emotional state (i.e. *Why* does this person feel an emotion?). The complete instruction is presented in the Supplementary Materials.

Stimuli for the OF-task were selected from the International Affective Picture System database (IAPS; Lang et al., 2008), the image set developed by the Kveraga lab (<http://www.kveragalab.org/stimuli.html>; Kveraga et al., 2015) and the internet (Google images). We selected images based on a pilot study, performed on an independent sample of subjects ( $n=22$ ). Details about this pilot study are available on request.

The OF-task was presented using the software package Presentation. The task presented thirty images on a black background in blocked fashion, with each block starting with a what, why or how cue (see Figure 1). Each image was shown three times, once for each cue type. Images were presented in blocks of six, each lasting 6 s, followed by a fixed inter trial interval of 2 s. Null-trials were inserted at random positions within the blocks. Both the order of the blocks and the specific stimuli within and across blocks were fully randomized, with a different randomization for each subject.

### Procedure

Each experimental session lasted about 2 h. Subjects who underwent two sessions had them on different days within a time-span of 1 week. On arrival, subjects gave informed consent and received thorough task instructions, including practice trials (see the Supplementary Materials for a translation of the task instructions). The actual time in the scanner was 55 min, and included a rough 3D scout image, shimming




Fig. 2. Schematic overview of the cross-validation procedures. (A) The partitioning of the dataset into an optimization-set (used for tuning of preprocessing and MVPA hyperparameters) and a validation-set (used to get a fully cross-validated, unbiased estimate of classification performance). The preprocessing and MVPA hyperparameters yielded from the optimization procedure were subsequently applied to the preprocessing and MVPA pipeline of the validation-set. (B) The within-subject MVPA pipeline of the self- and cross-analysis implemented in a repeated random subsampling scheme with 100 000 iterations. In each iteration, 90% of the self-data trials (i.e. train-set) were used for estimating the scaling parameters, performing feature selection and fitting the SVM. These steps of the pipeline (i.e. scaling, feature selection, SVM fitting) were subsequently applied to the independent test-set of both the self-data trials and the other-data trials.

sequence, 3-min structural T1-weighted scan, one functional run for the OF-task and two functional runs for the SF-task. We deliberately chose to present the SF-task after the OF-task to exclude the possibility that the SF-task affected the OF-task, thereby influencing the success of the decoding procedure.

After each scanning session, subjects rated their success rate for the SF-task and OF-task (see Supplementary Figure S1). In the second session, subjects filled out three personality questionnaires that will not be further discussed in this paper and were debriefed about the purpose of the study.

### Image acquisition

Subjects were tested using a Philips Achieva 3T MRI scanner and a 32-channel SENSE headcoil. A survey scan was made for spatial planning of the subsequent scans. Following the survey scan, a 3-min structural T1-weighted scan was acquired using 3D fast field echo (TR: 82 ms, TE: 38 ms, flip angle: 8°, FOV: 240 × 188 mm, 220 slices acquired using single-shot ascending slice order and a voxel size of 1.0 × 1.0 × 1.0 mm). After the T1-weighted scan, functional T2\*-weighted sequences were acquired using single shot gradient echo, echo planar imaging (TR = 2000 ms, TE = 27.63 ms, flip angle: 76.1°, FOV: 240 × 240 mm, in-plane resolution 64 × 64, 37 slices (with ascending acquisition), slice thickness 3 mm, slice gap 0.3 mm, voxel size 3 × 3 × 3 mm), covering the entire brain. For the SF-task, 301 volumes were acquired; for the OF-task 523 volumes were acquired.

### Model optimization procedure

As MVPA is a fairly novel technique, no consistent, optimal MVPA pipeline has been established (Etzel et al., 2011). Therefore, we adopted a validation strategy in the present study that is advised in the pattern classification field (Kay et al., 2008; Kriegeskorte et al., 2009). This strategy entailed that we separated our data into an optimization dataset to find the most optimal parameters for preprocessing and analysis, and a

validation dataset to independently verify classification success with those optimal parameters. We generated an optimization and validation dataset by running the SF-task and OF-task twice, in two identical experimental sessions for a set of thirteen subjects. The sessions were equally split between the optimization and validation set (see Figure 2A); first and second sessions were counterbalanced between the two sets. Based on a request received during the review process, we added nine new subjects to the validation dataset. Ultimately, the optimization-set held 13 sessions and the validation-set, after exclusion of 2 subjects (see Subjects section), held 20 sessions.

In the optimization-set, we explored how different preprocessing options and the so-called ‘hyperparameters’ in the MVPA pipeline affected the performance of the (multivariate) analyses (visualized in Figure 2B; see ‘MVPA pipeline’ subsection for more details). Thus, we performed the self- and cross-analyses on the data of the optimization set multiple times with different preprocessing options (i.e. smoothing kernel, low-pass filter and ICA-based denoising strategies) and MVPA hyperparameter values (i.e. univariate feature selection threshold and train/test size ratio during cross-validation). We determined the optimal parameters on the basis of classification performance, which was operationalized as the mean precision value after a repeated random subsampling procedure with 1000 iterations (see MVPA subsection for more details). A list with the results from the optimization procedure can be found in Supplementary Table S2 and Figure S2. The optimal parameters were then used for preprocessing and the self- and cross-analysis within the validation-set, in which the findings from the optimization-set were replicated. All findings discussed in the Results section follow from the validation-set (see Supplementary Figure S3 for an overview of the findings from the optimization-set).

### Preprocessing and single-trial modeling

Functional and structural data were preprocessed and analyzed using FSL 5.0 (Jenkinson et al., 2012) and MATLAB (2012b; www.



mathworks.com/products/matlab), using an in-house developed preprocessing pipeline and the parameters established in the optimization procedure. Functional data were corrected for motion (using FSL MCFLIRT) and slice timing and was spatially smoothed (5 mm isotropic kernel). After preprocessing, individual time series were modeled using a double  $\gamma$  hemodynamic response function in a single-trial GLM design using FSL's FEAT. Resulting  $\beta$ -values were converted to t-values (Misaki et al., 2010), constituting a whole-brain pattern of t-values per trial. Subsequently, the data were indexed by a gray-matter mask (excluding most white-matter, CSF and brainstem voxels). Thus, the data points for the MVPA consist of whole-brain (gray matter) t-value patterns per trial. For the optimization analyses, the data were transformed to standard space (MNI152, 2 mm) using FSL's FNIRT. To reduce computation time for the validation data, and in particular its corresponding permutation analysis, analyses on the validation dataset were performed on data in native (functional) space.

### Multi-voxel pattern analysis

**MVPA pipeline.** Within the optimization and validation dataset, we implemented an iterated cross-validation scheme that separated the data into a train-set and a test-set (this procedure is described in more detail in the next section). Before fitting the classifier on the train-set in each iteration of the cross-validation scheme, standardization and voxel selection were estimated and applied to the train-set. Standardization ensured that each feature (i.e. voxel) had zero mean and unit variance across trials. After standardization, voxel selection was performed in each iteration on the train-set by extracting the voxels with the highest average pairwise Euclidian distance across classes, which will be subsequently referred to as a voxel's differentiation score. More specifically, differentiation scores were calculated by subtracting the mean value across trials per class from each other (i.e. action—interoception, action—situation, interoception—situation), normalizing these values across voxels (yielding 'z-scores'), and taking their absolute value. The three resulting values per voxel were averaged and the most differentiating voxels (z-score threshold: 2.3, as determined by the optimization procedure; see Model optimization procedure section) were extracted and used as features when fitting the classifier. Importantly, the standardization parameters (voxel mean and variance) and voxel indices (i.e. which voxels had differentiation scores above threshold) were estimated from the train-set only and subsequently applied to the test-set to ensure independence between the train- and test-set (see Figure 2B). After standardization and voxel selection in each iteration, a support vector classifier (SVC) was fit on the train-set and cross-validated on the test-set, generating a class probability for each trial in the test-set. Our classifier of choice was the SVC implementation from the scikit-learn 'svm' module (Pedregosa et al., 2011) with a linear kernel, fixed regularization parameter (C) of 1.0, one-vs-one multiclass strategy, estimation of class probability output (instead of discrete class prediction) and otherwise default parameters.

**Cross-validation scheme and bagging procedure.** Cross-validation of the classification analysis was implemented using a repeated random subsampling cross-validation scheme (also known as Monte Carlo cross-validation), meaning that, for each iteration of the analysis, the classification pipeline (i.e. standardization, voxel selection and SVM fitting) was applied on a random subset of data points (i.e. the train-set) and cross-validated on the

remaining data (i.e. the test-set). Each trial belonged to one out of three classes: action, interoception or situation. Following the results from the parameter optimization process, we selected four trials per class for testing, amounting to 12 test-trials per iteration.

Per iteration, the classifier was fit on the train-set from the SF-data. Subsequently, this classifier was cross-validated on 12 test SF-trials (test-set 'self-analysis') and 12 test OF-trials (test-set 'cross-analysis'; see Figure 2B). This process was subsequently iterated 100 000 times to generate a set of class distributions for each trial. After all iterations, the final predicted class of each trial was determined by its highest summed class probability across iterations (also known as 'soft voting'; see Supplementary Figure S4). This strategy of a random subsampling cross-validation scheme in combination with majority (soft) voting is more commonly known as 'bagging' (Breiman, 1996). An important advantage of bagging is that it reduces model overfitting by averaging over an ensemble of models, which is especially useful for multi-voxel pattern analyses because fMRI data is known to display high variance (Varoquaux et al., preprint).

After generating a final prediction for all trials using the soft voting method, we constructed confusion matrices for both the self- and cross-analysis. In each raw confusion matrix with prediction counts per class, cells were normalized by dividing prediction counts by the sum over rows (i.e. the total amount of predictions per class), yielding precision-scores (also known as positive predictive value). In other words, this metric represents the ratio of true positives to the sum of true positives and false positives (see Supplementary Figure S5 for a description of the results expressed as recall estimates, or the ratio of true positives to the total number of samples in that class). This classification pipeline generated subject-specific confusion matrices that were subsequently averaged to generate the final classification scores.

**Statistical evaluation.** To evaluate the statistical significance of the observed average precision-scores in the confusion matrices, we permuted the original self- and cross-analysis 1300 times per subject with randomly shuffled class labels, yielding 1300 confusion matrices (with precision-scores). We then averaged the confusion matrices across subjects, yielding 1300 permuted confusion matrices reflecting the null-distribution of each cell of the matrix (which is centered around chance level classification, i.e. 33%). For each cell in the diagonal of the observed confusion matrix, P-values were calculated as the proportion of instances of values in the permuted matrix which were higher than the values in the observed matrix (Nichols and Holmes, 2002). To correct for multiple comparisons, P-values were tested against a Bonferroni-corrected threshold. The distribution of precision-scores and the relationship between precision-scores in the self- and cross-analysis is reported in Supplementary Figure S6.

**Spatial representation.** To visualize the classifier feature weights, we plotted the absolute feature weights averaged over iterations, subjects and pairwise classifiers (action vs interoception, action vs situation, interoception vs situation) that underlie our multiclass classification analysis. We chose to visualize the spatial representation of our model by plotting the average absolute feature weights, because the absolute value of feature weights in linear SVMs can be interpreted as how important the weights are in constructing the model's decision hyperplane (Guyon et al., 2002; Ethofer et al., 2009; Stelzer et al., 2014). To correct for




Fig. 3. Confusion matrices for the self- (left diagram) and cross-analysis (right diagram). Values indicate precision-scores, representing the proportion of true positives given all predictions for a certain class. Note that action and interoception columns in the cross-analysis confusion matrix do not add up to 1, which is caused by the fact that, for some subjects, no trials were predicted as action or interoception, rendering the calculation of precision ill-defined (i.e. division by zero). In this case, precision scores were set to zero.

a positive bias in plotting *absolute* weights, we ran the main classification analysis again with permuted labels to extract the average absolute feature weights that one would expect by chance. Subsequently, a voxel-wise independent t-test was performed for all feature weights across subjects, using the average permuted feature weights as the null-hypothesis, yielding an interpretable t-value map (see the supplementary code notebook on our Github repository for computational details).

### Additional analyses

In addition to the self-analysis and the self-to-other cross-analysis presented in the main text, we also performed a within-subjects other-to-self cross-analysis (see for a similar approach Corradi-Dell'Acqua et al., 2016) and a between-subjects self-analysis and self-to-other cross-analysis. These analyses forward largely similar results as the analyses presented in the main text. Due to space constraints, we present these additional analyses in the Supplementary Materials. Supplementary Figure S7 represents confusion matrices with precision and recall estimates for the other-to-self cross-analysis. Supplementary Figure S8 presents the results of MVPA analyses using condition-average voxel patterns across subjects instead of single-trial patterns within subjects.

### Univariate analysis

To be complete, we also report a set of univariate analyses performed on the SF-task and the OF-task data. The univariate analyses were performed on the validation dataset, and were subject to the same preprocessing steps as the MVPA analysis, except that we did not model each trial, but each condition as a separate regressor. The group-level analysis was performed with FSL's FLAME1 option. To examine differences in neural activity between conditions, we calculated contrasts between the three classes in the SF-task (self-action vs self-interoception; self-action vs self-situation and self-interoception vs self-situation) and the three classes in the OF-task (other-action vs other-interoception; other-action vs other-situation and other-interoception vs other-situation). We report clusters that were corrected using cluster-correction with a voxel-wise threshold of 0.005 ( $z = 2.7$ ) and a cluster-wise threshold of 0.05.

### Code availability

The MVPA-analysis and subsequent (statistical) analyses were implemented using custom Python scripts, which depend heavily on the *skbold* package, a set of tools for machine learning analyses of fMRI data developed in-house (see <https://github.com/lukassnoek/skbold>). The original scripts were documented and are hosted at the following Github repository: <https://github.com/lukassnoek/SharedStates>.

## Results

### Multi-voxel pattern analysis

The analyses of the SF-task demonstrated that voxel patterns reflecting imagined self-focused actions, interoceptive sensations and situations associated with emotion could be decoded accurately for each individual class (all  $P < 0.001$ , see Figure 3). Furthermore, when we generalized the classifier based on the SF-task to the data from the OF-task (i.e. cross-analysis), we found that neural representations of emotional actions, interoceptive sensations and situations of others could also be reliably decoded above chance (all  $P < 0.001$ ; see Figure 3). Supplementary Table S3 presents mean precision-scores across classes for each subject separately. As predicted, our findings demonstrate that *self-imagined* actions, interoceptive sensations and situations are associated with distinct neural patterns. Furthermore, and as predicted, our findings demonstrate that the patterns associated with *self-imagined* actions, sensations and situations can be used to decode *other-focused* actions, interoceptive sensations and situations (see Supplementary Figure S7 for the complementary other-to-self cross-analysis).

To visualize which neural regions were involved in the successful decoding of the three classes in the OF-task and SF-task, we display in Figure 4 the averaged absolute values of the SVM feature weights. Note that Figure 4 only displays one feature map, as both the self and cross-analysis depend on the same model. Regions displaying high and consistent feature weights across subjects were frontal pole (including parts of the dorsomedial prefrontal cortex and ventromedial prefrontal cortex), orbitofrontal cortex (OFC), inferior frontal gyrus (IFG), superior frontal gyrus (SFG), middle frontal gyrus (MFG), insular cortex, precentral gyrus, postcentral gyrus, posterior cingulate cortex/precuneus, superior parietal lobule (SPL), supramarginal gyrus (SMG), angular gyrus (AG), middle temporal gyrus (MTG), temporal pole (TP), lateral occipital cortex (LOC) and occipital pole




Fig. 4. Uncorrected t-value map of average feature weights across subjects; t-values were calculated by dividing the average absolute feature weights, which was corrected for positive bias by subtracting the mean permuted absolute weight across all iterations, by the standard error across subjects. Only voxels belonging to clusters of 20 or more voxels are shown.

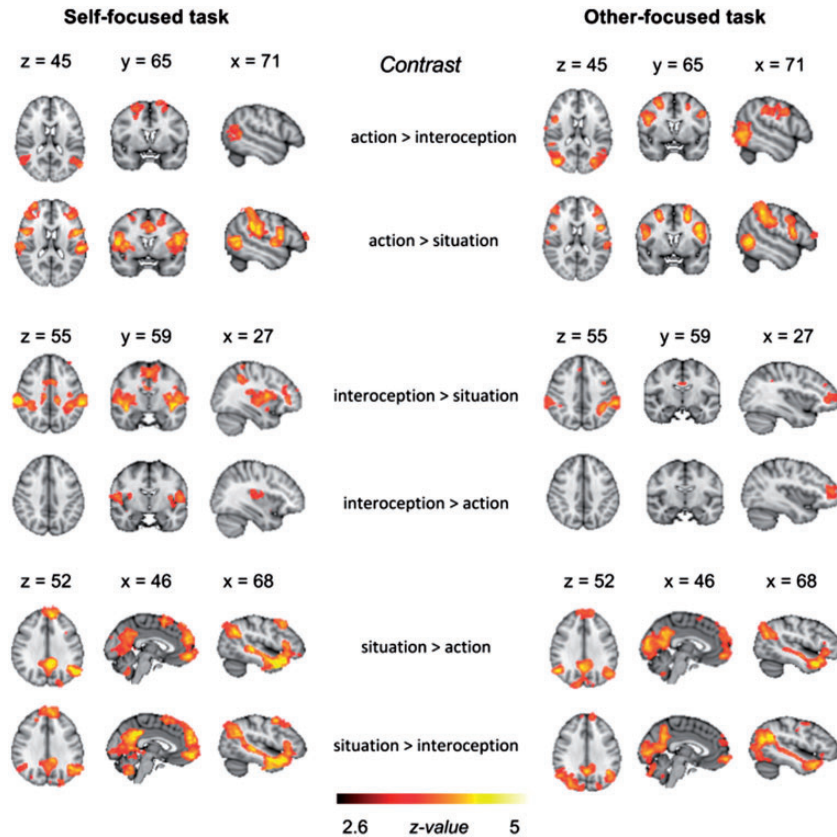


Fig. 5. Univariate contrasts for the self-focused and other-focused task.

(see Supplementary Table S4 for an overview of all involved regions).

### Univariate analyses

Figure 5 displays the pattern of neural activity revealed by univariate contrasts between the three different classes in the SF-task and the OF-task. For the sake of brevity, we summarize the most relevant univariate results here. Please see the Supplementary Materials and the study's Github repository for an overview of all clusters.

In the SF-task, action was associated with increased involvement of the MFG, SFG, AG, SMG, LOC and middle temporal gyrus (temporo-occipital) when compared with interoception, and increased involvement of the IFG, MFG, SFG, anterior cingulate cortex (ACC), supplementary motor area (SMA), precentral gyrus, postcentral gyrus, insular cortex, SMG, SPL, LOC and middle temporal gyrus (temporo-occipital) when compared with situation. Interoception was associated with increased

involvement of the insular cortex, precentral gyrus, postcentral gyrus and central operculum when compared with action, and increased involvement of the insular cortex, central operculum, parietal operculum, IFG, frontal pole, ACC, SMA, precentral gyrus, postcentral gyrus, SMG, SPL and putamen when compared with situation. The situation vs action contrast and the situation vs interoception contrast forwarded clusters in similar regions, including the temporal pole, superior/middle temporal gyrus, IFG, SFG, frontal pole, medial prefrontal cortex (mPFC), OFC, precuneus, posterior cingulate cortex (PCC), LOC, fusiform gyrus, hippocampus and lingual gyrus.

In the OF-task, action was associated with increased involvement of the IFG, MFG, SFG, precentral gyrus, postcentral gyrus, SMG, SPL, middle/inferior temporal gyrus (temporo-occipital), LOC and fusiform gyrus, when compared with interoception, and increased involvement of the IFG, MFG, SFG, frontal pole, precentral gyrus, postcentral gyrus, SMG, SPL, middle/inferior temporal gyrus (temporo-occipital) and LOC, when compared with situation. Interoception was associated with



increased involvement of the left frontal pole when compared with action, and increased involvement of the SMG, SPL, precentral gyrus, postcentral gyrus, PCC, IFG and frontal pole, when compared with situation. The situation vs action contrast and the situation vs interoception contrast forwarded clusters in similar regions, including the temporal pole, superior/middle temporal gyrus, frontal pole, mPFC, PCC, precuneus, AG, IOC, occipital pole, fusiform gyrus and lingual gyrus.

## Discussion

In this study, we investigated the neural overlap between self-focused emotion imagery and other-focused emotion understanding using a decoding approach. The results confirmed our hypothesis that other-focused representations of emotion-related actions, bodily sensations and situations can be decoded from neural patterns associated with accessing similar sources of information in a self-focused task. This cross-classification was successful even though the tasks employed different stimulus materials and instructions. Thus, the observed neural overlap between the underlying processes in the SF-task and OF-task cannot be attributed to similarities in stimulus dimensions or task instructions. Rather, we conclude from our findings that emotion experience and emotion understanding have basic psychological processes in common.

Although we could successfully classify the interoception class in the SF-task (across both datasets), and in the OF-task in the validation dataset, we were not able to successfully classify the interoception class in the OF-task in the optimization dataset. Furthermore, although precision and recall metrics demonstrated similar results for the action and situation cross-classification in the validation dataset, these metrics demonstrated different results for the classification of the interoception class (see Supplementary Figure S5). This difference was partly driven by the fact that trials were very infrequently classified as interoception in the cross-classification analysis. The finding that subjects reported lower success rates for the *what* trials in which they were asked to identify interoceptive sensations in other people than for the *how* (action) and *why* (situation) trials may point to a possible explanation for the inconsistent findings regarding interoception. Although speculative, it may be relatively easy to recognize (and represent) interoceptive sensations when they are described in words (as in the SF-task), but relatively hard to deduce these sensations when only diffuse cues about someone's internal state are available (e.g. posture, frowning facial expression, as in the OF-task).

An exploration of the spatial characteristics of the distributed neural pattern associated with successful decoding of the SF-task and OF-task revealed regions that are commonly active during self- and other-focused processing. First, we found that successful classification was associated with voxels in the precentral gyrus, IFG, SMA and SPL. These same regions were also revealed by the univariate analyses, in particular for the action and interoception classes. These regions are part of the so-called 'mirror' network, which is argued to support both action planning and action understanding (Gallese et al., 2004; Bastiaansen et al., 2009; Van Overwalle and Baetens, 2009; Spunt and Lieberman, 2013). Furthermore, we found that successful classification was associated with voxels in the lateral occipital cortex and fusiform gyrus, which have been linked in the literature to the processing of both concrete and abstract action (Wurm et al., 2015) and the (visual) processing of emotional scenes, faces and bodies (De Gelder et al., 2010; Sabatinelli et al.,

2011). The univariate analyses demonstrated activity in the IOC and the fusiform gyrus in particular for the situation class, both when subjects viewed images of other people in emotional situations, and when subjects imagined being in an emotional situation themselves.

Second, we found that successful classification was associated with voxels in regions associated with somatosensory processing (postcentral gyrus) and the representation of interoceptive sensations (insular cortex, see Craig, 2009; Medford and Critchley, 2010). Univariate analyses of the SF-task also demonstrated involvement of these regions for both the action and interoception classes. This pattern of activation is consistent with embodied cognition views that propose that thinking about or imagining bodily states is grounded in simulations of somatosensory and interoceptive sensations (Barsalou, 2009). In contrast to previous work on interoceptive simulation when observing pain or disgust in other people (cf. Bastiaansen et al., 2009; Lamm et al., 2011), the univariate analyses of the OF-task did not demonstrate insular cortex activation for the interoception class.

And third, we found that successful classification was associated with voxels in the middle temporal gyrus (including the temporal pole), PCC/precuneus, dmPFC and vmPFC. These regions are part of the so-called 'mentalizing' network (or 'default' network). This same network was also revealed by the univariate analyses, in particular for the situation class. Meta-analyses have demonstrated that the mentalizing network is commonly active during tasks involving emotion experience and perception (Lindquist et al., 2012), mentalizing/theory of mind (Spreng et al., 2009; Van Overwalle and Baetens, 2009), judgments about the self and others (Denny et al., 2012) and semantic/conceptual processing in general (Binder et al., 2009). Moreover, this network contributes to the representation of emotion knowledge (Peelen et al., 2010) and is involved in both empathy (Zaki and Ochsner, 2012; Keysers and Gazzola, 2014) and self-generated thought (Andrews-Hanna et al., 2014). We propose that this network supports the implementation of situated knowledge and personal experience that is necessary to generate rich mental models of emotional situations, both when experienced individually, and when understood in someone else (cf. Barrett and Satpute, 2013; Oosterwijk and Barrett, 2014).

The most important contribution of our study is that it provides direct evidence for the idea of shared neural resources between self- and other focused processes. It is important, however, to specify what we think this 'sharedness' entails. In research on pain, there is an ongoing discussion about whether experiencing pain and observing pain in others are distinct processes (Krishnan et al., 2016), or whether experiencing and observing pain involve a shared domain-specific representation (e.g. a discrete pain-specific brain state; Corradi-Dell'Acqua et al., 2016) and/or the sharing of domain-general processes (e.g. general negative affect; Zaki et al., 2016). Connecting to this discussion, we think that it is unlikely that our decoding success reflects the sharing of discrete experiential states between the SF-task and OF-task. After all, unlike in studies on pain, the stimuli in our tasks referred to a large variety of different actions, sensations and situations. Instead, decoding success in our study is most likely due to shared brain state configurations, reflecting the similar engagement of domain-general processes evoked by self- and other-focused instances of action (or interoceptive sensation or situation). This interpretation is consistent with views that suggests that global processes are shared between pain experience and pain observation (Lamm et al., 2011; Zaki et al., 2016) or between self- and other-focused tasks in



general (e.g. Legrand and Ruby, 2009). Moreover, this interpretation is consistent with the suggestion that neural re-use is a general principle of brain functioning (e.g. Anderson, 2016).

In our constructionist view, we posit that emotion imagery and understanding share basic psychological processes (cf. Oosterwijk and Barrett, 2014). More specifically, both emotion imagery and understanding are ‘conceptual acts’ in which the brain generates predictions based on concept knowledge (including sensorimotor and interoceptive predictions) that are meaningful within a particular situational context (see Barrett, 2012; Barrett and Simmons, 2015). Based on accumulating evidence, we propose that these predictions are implemented in domain-general brain networks (cf. Oosterwijk et al., 2012; Barrett and Satpute, 2013). The relative contribution of these networks depends on the demands of the situational context. Specifically, in contexts where people are focused on actions and expressions (their own or someone else’s) a network that supports the representation of sensorimotor states (i.e. the mirror system) may contribute relatively heavily; in contexts where people are focused on bodily states (their own or someone else’s) a network that supports the representation of interoceptive states (i.e. the salience network) may contribute relatively heavily; and in contexts where people are focused on interpreting a situation (their own or someone else’s) a network that supports a general inferential meaning-making function (i.e. the mentalizing network) may contribute relatively heavily (see also Oosterwijk et al., 2015). We believe that it is likely that our ability to successfully distinguish between classes in the self-task relies on the relatively *different* patterns of activity across these networks for actions, interoceptive sensations and situations. Regarding our ability to successfully generalize from the self- to the other-focused task, we believe that this relies on the relatively *similar* pattern of activity across these networks when people generate self-focused or other-focused instances of action (or interoceptive sensation or situation).

Our explicit manipulation of the weight of action, interoceptive and situational information in the SF-task and the OF-task tests the possibility of shared representation in a novel way. Although this procedure may seem artificial, social neuroscience studies support the notion that there is contextual variety in the contribution of action, interoceptive, and situation information when understanding other people (Van Overwalle and Baetens, 2009; Oosterwijk et al., 2015). Moreover, this weighting may mimic the variability with which these sources of information contribute to different instances of subjective emotional experience in reality (Barrett, 2012). In future directions, it may be relevant to apply the current paradigm to the study of individuals in which access to these sources of information is disturbed (e.g. individuals with different types of psychopathology) or facilitated (e.g. individuals with high interoceptive sensitivity).

In short, the present study demonstrates that the neural patterns that support imagining ‘performing an action’, ‘feeling a bodily sensation’ or ‘being in a situation’ are directly involved in understanding *other people’s* actions, sensations and situations. This supports our prediction that self- and other-focused emotion processes share resources in the brain.

## Acknowledgements

The authors would like to thank David Amodio for his helpful comments on a previous draft of this manuscript.

## Supplementary data

Supplementary data are available at SCAN online.

## Funding

The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Union Seventh Framework Programme (grant agreement no 275214) awarded to S.O. The writing of this paper was supported by a Netherlands Organization for Scientific Research VENI grant (451-13-008) awarded to S.O.

*Conflict of interest.* None declared.

## References

- Anderson, M. L. (2016). Précis of after phrenology: neural reuse and the interactive brain. *Behavioral and Brain Sciences*, **39**, e120.
- Andrews-Hanna, J.R., Smallwood, J., Spreng, R.N. (2014). The default network and self-generated thought: component processes, dynamic control, and clinical relevance. *Annals of the New York Academy of Sciences*, **1316**(1), 29–52.
- Barsalou, L.W. (2009). Simulation, situated conceptualization, and prediction. *Philosophical Transactions of the Royal Society of London: Biological Sciences*, **364**, 1281–9.
- Barrett, L.F. (2012). Emotions are real. *Emotion*, **12**, 413–29.
- Barrett, L.F., Satpute, A.B. (2013). Large-scale brain networks in affective and social neuroscience: towards and integrative functional architecture of the brain. *Current Opinion in Neurobiology*, **23**, 1–12.
- Barrett, L. F., Simmons, W. K. (2015). Interoceptive predictions in the brain. *Nature Reviews Neuroscience*, **16**, 419–29.
- Bastiaansen, J.A.C.J., Thouix, M., Keysers, C. (2009). Evidence for mirror systems in emotions. *Philosophical Transactions of the Royal Society B: Biological Science*, **364**, 2391–404.
- Binder, J.R., Desai, R.H., Graves, W.W., Conant, L.L. (2009). Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, **19**, 2767–96.
- Breiman, L. (1996). Bagging predictors. *Machine Learning*, **24**(2), 123–40.
- Brosch, T., Bar-David, E., Phelps, E.A. (2013). Implicit race bias decreases the similarity of neural representations of black and white faces. *Psychological Science*, **24**, 160–6.
- Carr, L., Iacoboni, M., Dubeau, M., Mazziotta, J.C., Lenzi, G.L. (2003). Neural mechanisms of empathy in humans: a relay from neural systems for imitation to limbic areas. *Proceedings of the National Academy of Sciences*, **100**, 5497–502.
- Corradi-Dell’Acqua, C., Tusche, A., Vuilleumier, P., Singer, T. (2016). Cross-modal representations of first-hand and vicarious pain, disgust and fairness in insular and cingulate cortex. *Nature Communications*, **7**, article number: 10904.
- Craig, A.D. (2009). How do you feel—now? The anterior insula and human awareness. *Nature Reviews Neuroscience*, **10**, 59–70.
- Decety, J. (2011). Dissecting the neural mechanisms mediating empathy. *Emotion Review*, **3**, 92–108.
- de Gelder, B., Van den Stock, J., Meeren, H.K., Sinke, C.B., Kret, M.E., Tamietto, M. (2010). Standing up for the body. Recent progress in uncovering the networks involved in the perception of bodies and bodily expressions. *Neuroscience & Biobehavioral Reviews*, **34**, 513–27.
- Denny, B.T., Kober, H., Wager, T.D., Ochsner, K.N. (2012). A meta-analysis of functional neuroimaging studies of self- and other

- judgments reveals a spatial gradient for mentalizing in medial prefrontal cortex. *Journal of Cognitive Neuroscience*, **24**, 1742–52.
- Ethofer, T., Van De Ville, D., Scherer, K., Vuilleumier, P. (2009). Decoding of emotional information in voice-sensitive cortices. *Current Biology*, **19**, 1028–33.
- Etzel, J.A., Valchev, N., Keysers, C. (2011). The impact of certain methodological choices on multivariate analysis of fMRI data with support vector machines. *Neuroimage*, **54**, 1159–67.
- Gallese, V., Keysers, C., Rizzolatti, G. (2004). A unifying view on the basis of social cognition. *Trends in Cognitive Sciences*, **8**, 396–403.
- Gilbert, S.J., Swencionis, J.K., Amodio, D.M. (2012). Evaluative vs. trait representation in intergroup social judgments: distinct roles of anterior temporal lobe and prefrontal cortex. *Neuropsychologia*, **50**(14), 3600–11.
- Guyon, I., Weston, J., Barnhill, S., Vapnik, V. (2002). Gene selection for cancer classification using support vector machines. *Machine Learning*, **46**(1–3), 389–422.
- Haynes, J.D. (2015). A primer on pattern-based approaches to fMRI: principles, pitfalls, and perspectives. *Neuron*, **87**(2), 257–70.
- Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M. (2012). FSL. *Neuroimage*, **62**, 782–90.
- Kay, K.N., Naselaris, T., Prenger, R.J., Gallant, J.L. (2008). Identifying natural images from human brain activity. *Nature*, **452**, 352–5.
- Keysers, C., Gazzola, V. (2014). Dissociating the ability and propensity for empathy. *Trends in Cognitive Sciences*, **18**, 163–6.
- Kriegeskorte, N., Simmons, W.K., Bellgowan, P.S., Baker, C.I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. *Nature Neuroscience*, **12**, 535–40.
- Krishnan, A., Woo, C.W., Chang, L.J., et al. (2016). Somatic and vicarious pain are represented by dissociable multivariate brain patterns. *Elife*, **5**, e15166.
- Kveraga, K., Boshyan, J., Adams, R. B., et al. (2015). If it bleeds, it leads: separating threat from mere negativity. *Social Cognitive and Affective Neuroscience*, **10**, 28–35.
- Lamm, C., Decety, J., Singer, T. (2011). Meta-analytic evidence for common and distinct neural networks associated with directly experienced pain and empathy for pain. *Neuroimage*, **54**, 2492–502.
- Lamm, C., Majdandžić, J. (2015). The role of shared neural activations, mirror neurons, and morality in empathy—a critical comment. *Neuroscience Research*, **90**, 15–24.
- Lang, P.J., Bradley, M.M., Cuthbert, B.N. (2008). *International Affective Picture System (IAPS): Affective Ratings of Pictures and Instruction Manual*. Technical Report A-8. Gainesville: University of Florida.
- Legrand, D., Ruby, P. (2009). What is self-specific? Theoretical investigation and critical review of neuroimaging results. *Psychological Review*, **116**, 252–82.
- Lench, H.C., Flores, S.A., Bench, S.W. (2011). Discrete emotions predict changes in cognition, judgment, experience, behavior, and physiology: a meta-analysis of experimental emotion elicitation. *Psychological Bulletin*, **137**, 834–55.
- Lindquist, K.A., Wager, T.D., Kober, H., Bliss-Moreau, E., Barrett, L.F. (2012). The brain basis of emotion: a meta-analytic review. *Behavioral Brain Sciences*, **35**, 121–43.
- Medford, N., Critchley, H.D. (2010). Conjoint activity of anterior insula and anterior cingulate cortex: awareness and response. *Brain Structure Function*, **214**, 535–49.
- Misaki, M., Kim, Y., Bandettini, P.A., Kriegeskorte, N. (2010). Comparison of multivariate classifiers and response normalizations for pattern-information fMRI. *Neuroimage*, **53**(1), 103–18.
- Mischkowski, D., Crocker, J., Way, B.M. (2016). From painkiller to empathy killer: acetaminophen (paracetamol) reduces empathy for pain. *Social Cognitive and Affective Neuroscience*, **11**, 1345–53.
- Nichols, T.E., Holmes, A.P. (2002). Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Human Brain Mapping*, **15**, 1–25.
- Norman, K.A., Polyn, S.M., Detre, G.J., Haxby, J.V. (2006). Beyond mind-reading: multi-voxel pattern analysis of FMRI data. *Trends in Cognitive Sciences*, **10**, 424–30.
- Oosterwijk, S., Barrett, L.F. (2014). Embodiment in the construction of emotion experience and emotion understanding. In: Shapiro, L., editor. *Routledge Handbook of Embodied Cognition*. New York: Routledge, 250–60.
- Oosterwijk, S., Lindquist, K.A., Anderson, E., Dautoff, R., Moriguchi, Y., Barrett, L.F. (2012). Emotions, body feelings, and thoughts share distributed neural networks. *Neuroimage*, **62**, 2110–28.
- Oosterwijk, S., Mackey, S., Winkielman, P., Wilson-Mendenhall, C., Paulus, M.P. (2015). Concepts in context: processing mental state concepts with internal or external focus involves different neural systems. *Social Neuroscience*, **10**, 294–307.
- Parkinson, C., Liu, S., Wheatley, T. (2014). A common cortical metric for spatial, temporal and social distance. *Journal of Neuroscience*, **34**, 1979–87.
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. (2011). Scikit-learn: machine learning in Python. *The Journal of Machine Learning Research*, **12**, 2825–30.
- Peelen, M.V., Atkinson, A.P., Vuilleumier, P. (2010). Supramodal representations of perceived emotions in the human brain. *Journal of Neuroscience*, **30**, 10127–34.
- Pulvermüller, F., Fadiga, L. (2010). Active perception: sensorimotor circuits as a cortical basis for language. *Nature Reviews Neuroscience*, **11**, 351–60.
- Rütgen, M., Seidel, E.M., Silani, G., et al. (2015). Placebo analgesia and its opioidergic regulation suggest that empathy for pain is grounded in self pain. *Proceedings of the National Academy of Sciences*, **112**, E5638–46.
- Sabatinelli, D., Fortune, E.E., Li, Q., et al. (2011). Emotional perception: meta-analyses of face and natural scene processing. *Neuroimage*, **54**, 2524–33.
- Singer, T. (2012). The past, present and future of social neuroscience: a European perspective. *Neuroimage*, **61**, 437–49.
- Spreng, R. N., Mar, R. A., Kim, A. S. (2009). The common neural basis of autobiographical memory, prospection, navigation, theory of mind, and the default mode: a quantitative meta-analysis. *Journal of Cognitive Neuroscience*, **21**, 489–510.
- Spunt, R.P., Lieberman, M.D. (2011). An integrative model of the neural systems supporting the comprehension of observed emotional behavior. *Neuroimage*, **59**, 3050–9.
- Spunt, R.P., Lieberman, M.D. (2013). The busy social brain: evidence for automaticity and control in the neural systems supporting social cognition and action understanding. *Psychological Science*, **24**, 80–6.
- Stelzer, J., Buschmann, T., Lohmann, G., Margulies, D.S., Trampel, R., Turner, R. (2014). Prioritizing spatial accuracy in high-resolution fMRI data using multivariate feature weight mapping. *Frontiers in Neuroscience*, **8**, 66.
- Uddin, L.Q., Iacoboni, M., Lange, C., Keenan, J.P. (2007). The self and social cognition: the role of cortical midline structures and mirror neurons. *Trends in Cognitive Sciences*, **11**, 153–7.

- Van Overwalle, F., Baetens, K. (2009). Understanding others' actions and goals by mirror and mentalizing systems: a meta-analysis. *Neuroimage*, **48**, 564–84.
- Varoquaux, G., Raamana, P.R., Engemann, D.A., Hoyos-Idobro, A., Schwartz, Y., Thirion, B. (2017). Assessing and tuning brain decoders: cross-validation, caveats, and guidelines. *NeuroImage*, **145**, 166–79.
- Waytz, A., Mitchell, J.P. (2011). The mechanisms for simulating other minds: dissociations between mirroring and self-projection. *Current Directions in Psychological Science*, **20**, 197–200.
- Wilson-Mendenhall, C.D., Barrett, L.F., Simmons, W.K., Barsalou, L. (2011). Grounding emotion in situated conceptualization. *Neuropsychologia*, **49**, 1105–27.
- Zaki, J., Ochsner, K.N. (2012). The neuroscience of empathy: progress, pitfalls and promise. *Nature Neuroscience*, **15**, 675–80.
- Zaki, J., Wager, T.D., Singer, T., Keysers, C., Gazzola, V. (2016). The anatomy of suffering: Understanding the relationship between nociceptive and empathic pain. *Trends in Cognitive Sciences*, **20**(4), 249–59.