



Maps of subjective feelings

Lauri Nummenmaa^{a,b,1}, Riitta Hari^{c,1}, Jari K. Hietanen^d, and Enrico Gleeran^{a,e,f,g}

^aTurku PET Centre and Turku University Hospital, University of Turku, FI-20520, Turku, Finland; ^bTurku University Hospital, University of Turku, FI-20520, Turku, Finland; ^cDepartment of Art, School of Arts, Design and Architecture, Aalto University, FI-00076, Espoo, Finland; ^dFaculty of Social Sciences and Psychology, University of Tampere, FI-33014, Tampere, Finland; ^eDepartment of Neuroscience and Biomedical Engineering, School of Science, Aalto University, FI-00076, Espoo, Finland; ^fDepartment of Computer Science, School of Science, Aalto University, FI-00076, Espoo, Finland; and ^gHelsinki Institute of Information Technology, Aalto University, FI-00076, Espoo, Finland

Contributed by Riitta Hari, July 3, 2018 (sent for review April 30, 2018; reviewed by Christian Keysers and Leonhard Schilbach)

Subjective feelings are a central feature of human life. We defined the organization and determinants of a feeling space involving 100 core feelings that ranged from cognitive and affective processes to somatic sensations and common illnesses. The feeling space was determined by a combination of basic dimension rating, similarity mapping, bodily sensation mapping, and neuroimaging meta-analysis. A total of 1,026 participants took part in online surveys where we assessed (i) for each feeling, the intensity of four hypothesized basic dimensions (mental experience, bodily sensation, emotion, and controllability), (ii) subjectively experienced similarity of the 100 feelings, and (iii) topography of bodily sensations associated with each feeling. Neural similarity between a subset of the feeling states was derived from the NeuroSynth meta-analysis database based on the data from 9,821 brain-imaging studies. All feelings were emotionally valenced and the saliency of bodily sensations correlated with the saliency of mental experiences associated with each feeling. Nonlinear dimensionality reduction revealed five feeling clusters: positive emotions, negative emotions, cognitive processes, somatic states and illnesses, and homeostatic states. Organization of the feeling space was best explained by basic dimensions of emotional valence, mental experiences, and bodily sensations. Subjectively felt similarity of feelings was associated with basic feeling dimensions and the topography of the corresponding bodily sensations. These findings reveal a map of subjective feelings that are categorical, emotional, and embodied.

feeling | emotion | consciousness | interoception | somatosensation

We humans constantly experience an ever-changing stream of subjective feelings that is only abolished during sleep, brain damage, and drug-induced altered states of the central nervous system. It is however most puzzling how our mental world is organized. External and internal pieces of information that proceed beyond nonconscious processing may be transformed into reliably reportable subjective experiences (1) or feelings that bear distinctive subjective qualia—the unique way how we feel different things. For example, trying to recall the title of the last book we read feels qualitatively different from the piercing headache during a high fever. The term feeling has multiple psychological and physiological definitions ranging from the subjectively accessible component of emotions to somatosensory experiences, ideas, and beliefs. Here, we use the word feeling to simply refer to the current, subjectively accessible phenomenological state of an individual. These inner sensations organize our mental lives and are also vital signs of our well-being (2, 3). Consequently, various unpleasant feelings are the most common reasons for seeking medical care (4). Despite the centrality of subjective feelings to ourselves and our conscious well-being, the relative organization and determinants of different feelings have remained poorly understood.

We experience in our bodies changes that relate to our well-being, physical activity, stress level, and emotional state. Thus, interoceptive and somatosensory inputs contribute significantly to the subjective feelings (5–7). Because the survival of an organism depends on its ability to maintain its physiology within an optimal homeostatic range, the ability to consciously monitor and feel certain physiological states (such as thirst and hunger) and detect

potential tissue damage has been critical already to our ancestors (3, 7). Such feelings vary markedly in their mental saliency. For example, heartbeat and digestive processes go most of the time unnoticed, whereas it is almost impossible to abolish the agony upon hearing of a loss of a dear friend. Such graded intensity of feelings allows filtering out unimportant ones from the limited-capacity working space in the executive system and focusing our executive problem-solving skills to deal with the most urgent ones (8).

Positive and negative emotions—physiological, motor, and cognitive programs—are inherently associated with subjective feelings (7, 9), and emotion is often considered as an essential building block of consciousness. Emergence of conscious feelings could thus require that the brain generates sensorimotor neural activation patterns of objects or events, as well as neural patterns regarding how (positively or negatively) these events change the state of the organism (10). Despite multiple classification schemes for specific emotions (11, 12), the contribution of emotion to other feeling states not traditionally considered as emotions (such as cognition and homeostatic processes), remains unknown. Finally, our ability to control different bodily and mental states varies greatly. A simple shift of attention across the room can be accomplished swiftly, while getting up from the bed at 4 AM to catch an early flight requires substantial effort, and chocking down the intense

Significance

Subjective feelings are a central feature of human life, yet their relative organization has remained elusive. We mapped the “human feeling space” for 100 core feelings ranging from cognitive and affective processes to somatic sensations; in the analysis, we combined basic dimension rating, similarity mapping, bodily sensation mapping, and neuroimaging meta-analysis. All feelings were emotionally loaded, and saliencies of bodily and mental experiences were correlated. Feelings formed five groups: positive emotions, negative emotions, cognitive processes, somatic states, and homeostatic states. Feeling space was best explained by emotionality, mental experience, and bodily sensation topographies. Subjectively felt similarity of feelings was associated with basic feeling dimensions and the bodily sensation maps. This shows that subjective feelings are categorical, emotional, and embodied.

Author contributions: L.N., R.H., J.K.H., and E.G. designed research; L.N. and E.G. performed research; L.N. and E.G. analyzed data; and L.N., R.H., J.K.H., and E.G. wrote the paper.

Reviewers: C.K., Netherlands Institute for Neuroscience; and L.S., Max Planck Institute of Psychiatry.

The authors declare no conflict of interest.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Data deposition: All datasets and resources are available at <https://doi.org/10.5281/zenodo.1291729>.

¹To whom correspondence may be addressed. Email: lauri.nummenmaa@utu.fi or riitta.hari@aalto.fi.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1807390115/-DCSupplemental.

grieving in one's partner's funeral would be next to impossible. The sense of controllability or agency—our experience of initiating, executing, and controlling thoughts and actions—is a central tenet of human phenomenological experience (13, 14). Because such sense of control binds our thoughts and actions to our ownership (13), it could also be one key dimension of subjective feelings.

Here, we address the fundamental principles that organize the human feeling states (*SI Appendix*, Fig. S1). We asked (*i*) how humans organize their feeling states, (*ii*) what kind of mental experiences and bodily sensations would best explain the representational structure of the feelings, and (*iii*) whether the mental experiences and bodily sensations are associated with distinct neural activation patterns. We focused on mapping the basic dimensions (Experiment 1), ontology (Experiment 2), as well as bodily (Experiment 3) and neural (meta-analysis and synthesis of Experiments 1–3) basis of a broad array of feeling states (*SI Appendix*, Fig. S2 and Table S1). We first quantified the relative intensities of four hypothesized core subjective dimensions (intensity of bodily sensations, saliency of mental experience, emotional valence, and agency) of 100 common subjective feelings spanning from homeostatic (e.g., hunger) and emotional (e.g., pleasure) states to cognitive functions (e.g., recalling). We also measured the relative frequency of experiencing each feeling as the lapse since the last remembered occurrence of each feeling. Next, we measured the experienced similarity of these subjective feelings and mapped the topography of bodily sensations associated with each feeling. Neural activation patterns associated with each state were derived using large-scale meta-analysis of fMRI data. We quantified the spatial representations of these states and linked the representational organization of the subjective states with their bodily and neural activation patterns. We show that subjective mental states are embodied and emotionally valenced, and that there is a clear correspondence between the mental experiences and their bodily basis that also pertains to the underlying neural activation patterns in the bodily domain.

Results

Experiment 1: Basic Dimensions of Subjective Feelings. Fig. 1 shows the associations between the core dimensions of subjective feelings, and *SI Appendix*, Fig. S2 shows the corresponding probability density functions of the data aggregated across subjects (see *SI Appendix*, Table S1 for means and SEs of mean). All tested states were associated with salient mental and bodily sensations. The relative strength of the mental experiences and bodily sensations differed greatly across states but were linearly associated, when fitted with separate regression lines for two major clusters in the data. The majority of the feeling states were experienced as emotional—either pleasant or unpleasant. Only few, mainly physiological and cognitive, states were experienced as affectively neutral. As expected, emotional intensity (i.e., deviation from the midpoint neutral affective state toward pleasure or displeasure) was significantly associated with the intensity of both mental and bodily sensations. Pleasant and controllable states were more frequent than unpleasant and uncontrollable ones. Finally, subjects felt less control over unpleasant than pleasant feelings and less control over bodily than mental states.

Experiment 2: Mapping the Mental Feeling Space. We next mapped the topographical organization of the feeling states based on the similarity ratings obtained in Experiment 2. Fig. 2 shows the 2D representation of the feeling space. The density-based clustering (DBSCAN; ref. 15) and t-distributed stochastic neighbor embedding (t-SNE; ref. 16) solutions revealed five discrete clusters: positive emotions, negative emotions, cognitive processes, somatic states and illnesses, and homeostatic states. The topographical dimension intensity heatmaps (based on Experiment 1) were then used to evaluate how the states in each cluster differ with respect to the underlying basic dimensions. The two dimensions of the t-SNE

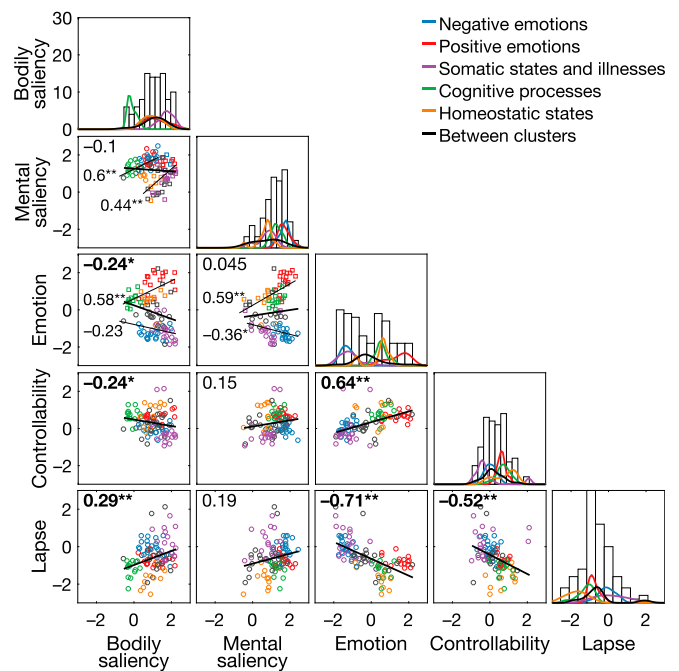


Fig. 1. Associations between the basic dimensions of subjective experience. Histograms show the distribution for each evaluated dimension with overlaid probability density curves per clustering in Experiment 2. Scatterplots show median z-scores for each of the 100 feelings (see *SI Appendix*, Fig. S2 for the complete distribution densities). Color coding of the tokens is based on clustering per Experiment 2. Normalized data are modeled with one or two LS regression lines depending on the optimal fit. All reported values are Spearman's correlations $*P < 0.05$, $**P < 0.005$.

derived feeling space mapped primarily to the strength of mental feelings on vertical axis and emotional valence (pleasant–unpleasant) on the horizontal axis, so that both positive and negative emotional states were experienced strongly, and homeostatic processes and illnesses weakly in the mind. Emotional valence was also associated with the sense of control, in that sensations associated with cognition and positive emotions were the easiest, and those associated with illnesses and negative emotions the hardest to control. Intensity of bodily sensations mapped to the bottom corners of the space. The feeling spaces were concordant across participants with a mean Spearman's correlation split-half reliability of 0.94 (*SI Appendix*, Fig. S3).

Fig. 3 shows the bodily sensations associated with each feeling state, revealing that almost all feeling states were associated with a unique, discernible bodily sensation map. Although the organization of these bodily maps was associated with the organization of the subjective organization of the feeling states (i.e., Experiment 2), t-SNE mapping based on the bodily sensations revealed more clearly a linear continuum along the distributed versus focal fingerprints of bodily feelings (*SI Appendix*, Figs. S4 and S5).

Representational Similarity Across Neural Activation and Domains of Subjective Feelings. Finally, we assessed the representational similarity between the measured features of subjective experience. This across-dimension representational similarity analysis (RSA; ref. 17) (Fig. 4) revealed that the ground-truth similarity metric (i.e., subjectively felt organization) was significantly associated with all of the basic dimensions, as well as with bodily sensations and semantic similarity. Bodily sensation maps were also significantly associated with all other dimensions except controllability and lapse. Semantic similarity was associated with bodily saliency, emotion, bodily sensation maps, and neural similarity (NeuroSynth). Neural similarity was, in turn, associated with similarity of bodily

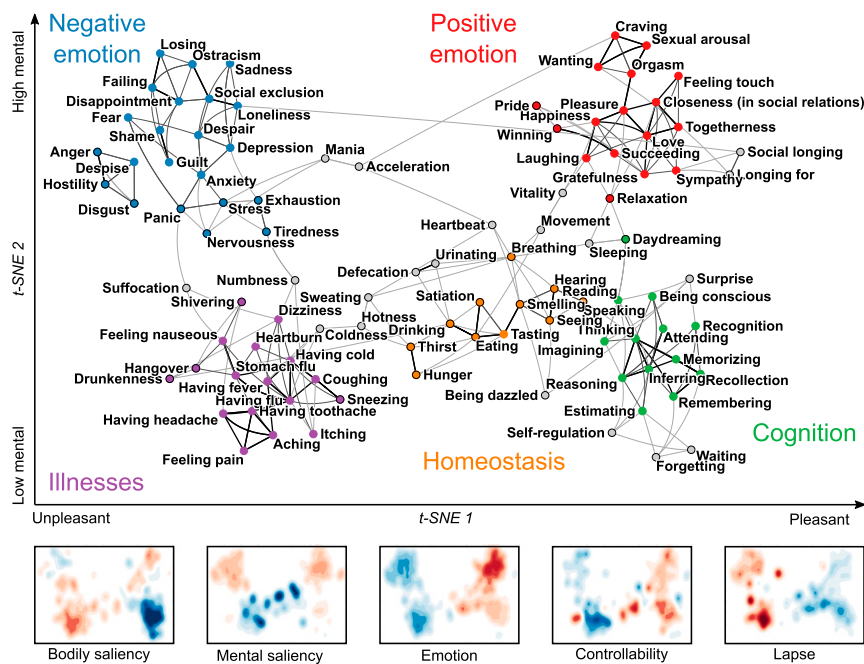


Fig. 2. The feeling space. (Upper) Two-dimensional map of the feeling space based on the sorting task average distance matrix between items arranged by t-SNE and clustered with DBSCAN. Color coding indicates cluster structure; gray feelings do not belong clearly to any cluster. Colored items with black edge are DBSCAN border elements. To retain the information of the distance matrix the closest three items for each node are connected with lines. Thick dark lines are showing distances that belong to the top 33rd percentile of the visualized lines (i.e., the closest items). (Lower) Heatmaps showing how strongly each basic dimension of subjective experience is associated with each discrete feeling at each location of the feeling space. Color coding shows the relative intensity as median z-score (as in Fig. 1) from high (red) to low (blue).

saliency, bodily sensation maps, and semantic similarity [all P values < 0.05 , false discovery rate (FDR) corrected].

Discussion

Our data provide a detailed map of the human feeling space, where subjective feelings were strongly coupled with bodily sensations, and nearly all subjective experiences were qualified by emotional tone. Subjective states were best described on a 2D map with five distinct feeling clusters. Representational similarity analysis revealed strong correspondence between the mental feeling space and the corresponding bodily sensations, basic dimensions of subjective experience, as well as similarity between neural and bodily sensation maps across different subjective feelings. Altogether these findings show that feeling states are categorical, emotional, and embodied.

Organization of the Feeling Space. Subjective feelings were organized as a 2D space with discrete five-cluster structure (positive emotions, negative emotions, cognitive processes, somatic states and illnesses, and homeostatic and sensory states). This mapping was concordant across individuals, suggesting that the feeling space is minimally influenced by individual differences. Thus, its organization likely reflects inherent coupling between the sensory and interoceptive systems and the frontoparietal brain circuit supporting consciously accessible feelings (18). When linked with the topography intensity of basic feeling-dimension ratings, the feeling space was best described by the orthogonal dimensions of affective valence (pleasant–unpleasant) and mental saliency (low–high), with bodily sensations being most profoundly experienced on the anti-diagonal (positive emotions and illnesses; Fig. 2). The clustering of the feelings likely reflects the way the brain binds and weighs the interoceptive and exteroceptive inputs for generating each feeling state.

When individuals evaluate the presence of mental functions in other agents, the presence of physiological, emotional, and cognitive-perceptual experiences, and the sense of agency, are the most important determinants guiding the decision (19, 20). These dimensions

were salient also in the present data, indicating that our own feelings provide the anchors against which to attribute mental states of other beings. Shared mapping of feelings across individuals may thus serve important social functions. For example, the ability to exchange conscious states across individuals via action–perception loops is an essential building block of sociability (21), and the capability to promote social cohesion by sharing feelings has likely yielded significant evolutionary advantages already to our ancestors (22).

Feelings Are Coupled with Bodily Processes. Feelings were systematically referenced to bodily states, even for states considered as purely cognitive, such as attending or reasoning. Additionally, the more strongly some feeling was experienced in the body, the more salient it was mentally. Different feelings were associated with distinct bodily fingerprints, and representational structure of these bodily maps converged with the mentally experienced structure of feelings: The qualia of the mental experience of any given feeling were associated with the unique topography of the concurrent bodily sensations. This finding accords with results showing that voluntary reproduction of bodily states associated with, for example, specific emotions leads to the subjective feeling of the corresponding emotion (23). Accordingly, an individual's current emotional feeling can be accurately classified using signals from the primary somatosensory and insular cortices (24). Thus, interoceptive signaling—associated primarily with homeostasis and a threatening tissue damage—is an important feature of subjective feelings in general (3, 7), similarly as has been previously established specifically for emotional states (25, 26). Such embodiment of subjective experiences may also extend to self-awareness: Adults and children consistently think of “self” as occupying a physical location within the body—oftentimes in head, but also remarkably often in the torso (27). However, even when some types of spinal-cord lesions prevent interoceptive signals from entering the brain, they do not necessarily abolish subjective awareness; for example, patients with pure autonomic failure can still experience embodied



Fig. 3. Bodily feeling maps. Pixel intensities show regions where each feeling was associated with statistically significant ($P < 0.05$, FDR corrected) bodily sensations. The data are arranged into the matrix approximately as in Fig. 2; clustering is per Experiment 2. Colorbar indicates the effect size. See *SI Appendix*, Figs. S4 and S5 for spatial arrangement per t-SNE.

subjective states, such as emotions (28). Moreover, complete bilateral lesions to the interoceptive regions in the insula do not abolish emotions (29). Bodily feedback cannot thus be the sole contributor to subjective feelings, as was also confirmed by our representational similarity analysis.

Emotion as a Core Dimension of Subjective Feeling. Nearly all subjective feelings were imbued with emotional qualities. During emotions, the benefit versus harm associated with external and internal events triggers an integrated blend of good or bad feelings, or vigorous activation or relaxation. Such affective sensations involve cognitive, perceptual, and physiological processes and thus constitute the core of human subjective experience in general (9). Emotional intensity emerged as the best predictor for the organization of the subjective feeling space in the RSA. Despite both scientists (30) and laypersons alike often considering affect and cognition at least partially independent, we found that most feelings actually carry affective valence, further underlining the importance of affect to conscious states in general (10). Prior studies have focused on mapping the feelings associated with prototypical emotions (11, 12). However, the present data show that canonical basic and social emotions occupy a limited area of the total feeling space, which covers a wide array of mental and bodily functions and states. This result fits with the reasoning that the same frontoparietal brain systems stand as the basis of both emotional and nonemotional feelings but so that input from subcortical and other cortical systems associated with the

physiological and behavioral aspects of emotions is needed to add the affective quality to certain feelings (18).

Links Among Mind, Body, and the Brain in Generating Feelings. The representational similarity analysis revealed that the subjective feelings have multiple determinants. Using the arrangement-based organization as the “ground truth” metric, we found that it was associated with both the organization of the feeling-specific bodily topographies as well as with the relative intensities of the basic dimensions of feelings. The organization of subjective feelings at mental and embodied level, however, corresponded with the pure semantic (“is a”) similarities only weakly, confirming that the feeling space does not simply reflect the interchangeability of the semantic concepts (i.e., names of the feelings) used in the study (31). Instead, each feeling state is likely characterized by a combination of interoceptive and exteroceptive inputs evaluated also with respect to emotional qualities. We propose that humans intuitively bind these noisy inputs into coherent subjective experiences, similarly as has been previously found to happen for sensory percepts, so that high pitches, bright colors, and high spatial positions go together (32).

Neural organization of feelings was also associated with their bodily basis and database-determined semantic similarity, providing evidence of specific neural signatures of specific subjectively felt bodily states (10, 26). However, such direct mapping does not extend to mental states, and no association was found between the subjective and neural dimensions. This finding is not fully




Fig. 4. Representational similarity between determinants of subjective experience (Spearman's correlation). Each cell in the matrix represents pairwise RSA between two similarity matrices for the 100 feelings; * $P < 0.05$ (FDR corrected), ** $P < 0.005$ (FDR corrected, permutation-based significance thresholding).

unexpected, given that gross neural activation patterns associated with, for example, seeing (occipital cortex) and hearing (temporal cortex) are spatially divergent, despite felt similarity of the processes. Thus, unlike bodily feeling states, higher-order subjective mental states do not necessarily reduce to local brain activations. However, our analysis is constrained by the spatial accuracy of fMRI and the univariate nature of the NeuroSynth database. Prior studies employing multivariate pattern recognition have indeed linked the similarity of neural activation patterns with the similarity of subjective sensory percepts (33) and feelings (24), suggesting a direct link between brain activity and subjective experiences.

Limitations. We used a relatively crude distinction of the subjective states—for example, “seeing” would likely feel very different when we are seeing a growling black bear rather than a cute baby. However, our goal was at this stage to characterize the relationship of the most fundamental aspects of human feelings, rather than the fine-grained dissection of the qualia within each possible domain. Explaining how the distinct qualia, such as the pleasure derived from having sex or the agony triggered by seeing our loved ones suffer, are associated with neuronal activity is known as the “hard problem” of consciousness (34). Although our data clearly point toward bodily and emotional basis of subjective feelings, this connection does not obviously explain how bodily sensations are transformed into subjective qualia. This problem is further highlighted by the lack of representational similarity between subjective feelings and their neural signatures (Fig. 4): We do not understand well how these feelings are generated in the brain. Finally, humans have a limited access to the workings of their higher-order cognitive processes (35), and we cannot ascertain how closely the self-reported subjective sensations associate with the underlying neural and cognitive machinery.

Conclusions

Why do we experience anything at all? The present data suggest that we feel bodily events that are potentially good or bad to us. We found that subjective feelings are categorical, emotionally valenced, and embodied, and that the valence and mental involvement are the most important dimensions of the subjective feelings. These results extend the current models of the role of somatosensation (36) and embodiment (26, 37) in emotional

processing to more general-level mechanisms determining the nature of human experiences.

Materials and Methods

The original list of the feeling tokens was derived from the American Psychological Association Glossary of Psychological Terms (<https://dictionary.apa.org/>), complemented with basic physiological and homeostatic processes, as well as with common illnesses. We then selected a total of 100 core feelings (SI Appendix, Table S1) belonging to seven broad a priori categories: cognition (e.g., thinking, reasoning), sensation and perception (e.g., seeing, hearing), motivational states (e.g., wanting, craving), emotional states (e.g., anger, fear), homeostatic states (e.g., hunger, thirst), physiological processes (e.g., breathing, sleeping), common illnesses (e.g., flu, fever), and feelings central to common psychiatric disorders (e.g., depression, anxiety). A pilot study confirmed that these sensations cover both frequently and infrequently experienced feeling states.

SI Appendix, Fig. S1 summarizes the experimental protocol. Altogether 1,026 participants (880 females; 146 males) completed the experiments (Experiment 1: $n = 339$, $MD_{age} = 27$ y, 300 females, 39 males; Experiment 2: $n = 226$, $MD_{age} = 27$ y, 181 females, 45 males; Experiment 3: $n = 461$, $MD_{age} = 26$ y, 399 females, 62 males). They were recruited from university email lists and social media. All experimental studies were conducted online. The participants first gave informed consent and provided background information (age, sex, and education). After this, the participants were provided with the instructions for the experiment. Data acquisition and analysis code and the datasets are available at <https://doi.org/10.5281/zenodo.1291729>. Aalto University Institutional Review Board approved all experiments involving human subjects.

Experiment 1: Basic Dimensions of Subjective Feelings. We hypothesized that four main dimensions (saliency of mental experience, saliency of bodily sensation, emotion, and controllability) would underlie the feeling space. To map these dimensions to specific feelings, the participants were shown the name of one candidate feeling at a time on the computer screen. For each feeling, the participant was asked to evaluate (i) how much that feeling is felt in the body (not at all—extremely much), (ii) how much that feeling is experienced in the mind (not at all—extremely much), (iii) how pleasant it feels (extremely unpleasant—extremely pleasant), and (iv) how much control they feel having on it (not at all—extremely much) on visual analog scales (VAS). An additional VAS was used for measuring (v) the frequency of experiencing the feelings, based on the last remembered occurrence of each feeling (within the last hour—last year, or not at all). Due to the large number of tokens (100), each participant rated 20 randomly chosen feeling tokens; they were asked to rate more tokens if they were willing to. Participants not completing at least one full batch (20) of tokens were excluded ($n = 88$). Associations between the basic dimensions were estimated using Pearson correlation and by fitting least-square regression lines between the dimensions. Because Experiment 2 revealed a clear cluster structure of the feelings, we fitted either one or two regression lines to each pair of variables depending on the optimal fit.

Experiment 2: Mapping the Mental Feeling Space. We next mapped the mental feeling space using similarity rating using a modified computer-based Q-sort task. Subjects were shown all tokens as a randomly ordered list on the left side of the screen, and they were asked to arrange them to a rectangular box on the right so that states feeling similar would be close to each other. This technique is known to yield consistent results with direct pairwise ratings ($r_{is} > 0.9$; ref. 38). It was stressed that the arrangement should be based on the similarity of feelings, rather than on conceptual or semantic relatedness. Participants could adjust the placement as long as they wanted. Participants moving less than 25% of the tokens ($n = 138$) were excluded from the data, and tokens not moved into the target area were not considered in the analyses. Pairwise similarities between the tokens were extracted as their Euclidian distance. Distances were scaled so that the maximum distance across all participants was set to 1. Consistency of participant-wise feeling space arrangement was estimated using split-half reliability (Spearman's correlation) over 5,000 random splits. Mean distance matrix was constructed (median distance matrix was essentially similar, $r = 0.99$ $P < 0.001$ in Mantel test) and subjected to DBSCAN (15) to obtain reproducible clusters as well as “border” elements, and “outlier” elements that do not fit in a single cluster. Next, t-SNE was used to visualize the underlying structure of the phenomenological space (16).

To test how different basic dimensions of feelings determined in Experiment 1 project to the subjective feeling space determined in Experiment 2, we modeled, for each data point in the t-SNE-derived feeling space, the intensity of each basic dimension with a Gaussian kernel with 8-unit FWHM weighted

with the mean basic dimension rating of that token in Experiment 2. The resulting maps thus reveal how strongly each basic dimension of subjective experience is felt across the 2D feeling space. This representation, based on t-SNE and the basic dimensions, thus merges the traditional network and space-based representations often employed in semantic analyses of concepts used in natural language (31), and it provides straightforward access to the organization of the subjective feelings at multiple levels of analysis.

Experiment 3: Mapping the Bodily Sensation Space. Bodily topography of feelings was mapped using the emBODY tool (ref. 26; <https://version.aalto.fi/gitlab/eglerean/embody>, ref. 39). Subjects were shown one candidate feeling at a time, and were asked to color, in a blank body shape, whereabouts in the body the state feels. Due to a large number of tokens, each participant completed the ratings for 20 randomly chosen tokens after which they were given an option to complete another batch of 20 until they quit or had completed all of the 100 tokens. Subjects not completing at least one whole batch of 10 tokens were excluded. The resulting subject-wise bodily sensation maps were averaged for each state and subjected to random effects analysis. Bodily sensation maps (BSMs) for each state were obtained by mass univariate *t* tests against zero feeling values (with FDR correction across all pixels and all tokens). Similarity matrices for the different feelings were constructed by computing Euclidean distance between the unthresholded mean BSM effect size maps between each pair of feelings.

Meta-Analysis of Neural Activation Patterns Across Feeling States. To estimate the feeling state-wise neural activation patterns, we employed data from Neurosynth database (www.neurosynth.org). This tool allows automated meta-analysis of brain activations from tens of thousands of fMRI studies (ref. 40; analysis date November 15, 2016). We first obtained the full list of subject terms of studies stored in Neurosynth. Two independent raters evaluated the best possible (if any) matches from the Neurosynth term list with the subjective states used in the behavioral experiments. Thresholded Neurosynth reverse inference maps (summarizing altogether 9,821 brain imaging experiments) were downloaded for the terms found in the database. If two Neurosynth

terms matched with one token (e.g., “attention” and “attending”), their maps were averaged. Altogether we found NeuroSynth database matches for 44 of the 100 tokens. Similarity between the meta-analytic activation maps was determined as the Spearman’s correlation between their voxels.

Representational Similarity Across Neural Activation and Domains of Subjective Feelings. We next asked how strongly the different dimensions of subjective feelings (Experiment 1—Basic dimensions; Experiment 2—Mental feeling space; Experiment 3—Bodily sensation space; Meta-analysis—Neural activation patterns) are associated with each other as well as with the brain activation patterns underlying each state. We also tested whether the similarity structure of feelings or their bodily and neural underpinnings would reflect only their semantic relatedness, rather than subjectively felt similarity of the feeling themselves. We thus computed the semantic distance between each word pair using the SimService tool (ref. 41; swoogle.umbc.edu/SimService/). Semantic distance was computed using both Latent Semantic Analysis (LSA) based on two corpora (Refined Stanford WebBase corpus, and LDC English Gigawords Corpus) and WordNet relationships between words. The tool provides four measures of similarity between word pairs, which were averaged to index how interchangeable two words are (31), providing a statistical estimate of the similarity of the semantic concepts referring to the feelings as they are used in language. To test for the representational similarity of feelings across different domains and dimensions, we used RSA (17) between the off-diagonal elements of the similarity matrixes derived from Experiments 1–3, LSA + WordNet, and NeuroSynth. This resulted in total nine similarity structures: Basic dimensions (Experiment 1: intensity of bodily feelings and mental involvement, affect, and controllability), frequency (Experiment 1), subjective feeling space (Experiment 2), bodily sensation maps (Experiment 3), and the data-driven neural (NeuroSynth) and semantic (LSA + WordNet) similarity structures.

ACKNOWLEDGMENTS. We acknowledge the computational resources provided by The Aalto Science-IT project. This research was supported by the Academy of Finland Grants 265917, 294897, and 304385 and European Research Council Starting Grant 313000 (to L.N.).

- Dehaene S, Changeux JP (2011) Experimental and theoretical approaches to conscious processing. *Neuron* 70:200–227.
- Fredrickson BL (2004) The broaden-and-build theory of positive emotions. *Philos Trans R Soc Lond B Biol Sci* 359:1367–1378.
- Damasio A (1999) *The Feeling of What Happens: Body and Emotion in the Making of Consciousness* (Harcourt Brace, New York).
- St Sauver JL, et al. (2013) Why patients visit their doctors: Assessing the most prevalent conditions in a defined American population. *Mayo Clin Proc* 88:56–67.
- Wilson AD, Golonka S (2013) Embodied cognition is not what you think it is. *Front Psychol* 4:58.
- Craig AD (2002) How do you feel? Interoception: The sense of the physiological condition of the body. *Nat Rev Neurosci* 3:655–666.
- Damasio A, Carvalho GB (2013) The nature of feelings: Evolutionary and neurobiological origins. *Nat Rev Neurosci* 14:143–152.
- Damasio A (1994) *Descartes’s Error: Emotion Reason and the Human Brain* (Grosset Putnam, New York).
- Russell JA (2003) Core affect and the psychological construction of emotion. *Psychol Rev* 110:145–172.
- Damasio A, Meyer K (2009) Consciousness: An overview of the phenomenon and of its possible neural basis. *The Neurology of Consciousness* (Academic, San Diego), pp 3–14.
- Toivonen R, et al. (2012) Networks of emotion concepts. *PLoS One* 7:e28883.
- Tellegen A, Watson D, Clark LA (1999) On the dimensional and hierarchical structure of affect. *Psychol Sci* 10:297–303.
- Gallagher I (2000) Philosophical conceptions of the self: Implications for cognitive science. *Trends Cogn Sci* 4:14–21.
- David N, Newen A, Vogeley K (2008) The “sense of agency” and its underlying cognitive and neural mechanisms. *Conscious Cogn* 17:523–534.
- Ester M, Kriegel H-P, Sander J, Xu X (1996) A density-based algorithm for discovering clusters a density-based algorithm for discovering clusters in large spatial databases with noise. *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining* (AAAI Press, Portland, OR), pp 226–231.
- van der Maaten L, Hinton G (2008) Visualizing data using t-SNE. *J Mach Learn Res* 9: 2579–2605.
- Kriegeskorte N, Mur M, Bandettini P (2008) Representational similarity analysis—Connecting the branches of systems neuroscience. *Front Syst Neurosci* 2:4.
- LeDoux JE, Brown R (2017) A higher-order theory of emotional consciousness. *Proc Natl Acad Sci USA* 114:E2016–E2025.
- Weisman K, Dweck CS, Markman EM (2017) Rethinking people’s conceptions of mental life. *Proc Natl Acad Sci USA* 114:11374–11379.
- Gray HM, Gray K, Wegner DM (2007) Dimensions of mind perception. *Science* 315:619.
- Hari R, Kujala MV (2009) Brain basis of human social interaction: From concepts to brain imaging. *Physiol Rev* 89:453–479.
- Dunbar RIM (1996) *Grooming, Gossip, and the Evolution of Language* (Harvard Univ Press, Cambridge, MA).
- Philippot P, Chapelle G, Blairy S (2002) Respiratory feedback in the generation of emotion. *Cogn Emot* 16:605–627.
- Saarmäki H, et al. (2016) Discrete neural signatures of basic emotions. *Cereb Cortex* 26:2563–2573.
- James W (1884) What is an emotion? *Mind* 9:188–205.
- Nummenmaa L, Glerean E, Hari R, Hietanen JK (2014) Bodily maps of emotions. *Proc Natl Acad Sci USA* 111:646–651.
- Starmans C, Bloom P (2012) Windows to the soul: Children and adults see the eyes as the location of the self. *Cognition* 123:313–318.
- Heims HC, Critchley HD, Dolan R, Mathias CJ, Cipolotti L (2004) Social and motivational functioning is not critically dependent on feedback of autonomic responses: Neuropsychological evidence from patients with pure autonomic failure. *Neuropsychologia* 42:1979–1988.
- Damasio A, Damasio H, Tranel D (2013) Persistence of feelings and sentience after bilateral damage of the insula. *Cereb Cortex* 23:833–846.
- Zajonc RB (1980) Feeling and thinking - PREFERENCES NEED no INFERENCEs. *Am Psychol* 35:151–175.
- Griffiths TL, Steyvers M, Tenenbaum JB (2007) Topics in semantic representation. *Psychol Rev* 114:211–244.
- Spence C (2011) Crossmodal correspondences: A tutorial review. *Atten Percept Psychophys* 73:971–995.
- Said CP, Moore CD, Engell AD, Todorov A, Haxby JV (2010) Distributed representations of dynamic facial expressions in the superior temporal sulcus. *J Vis* 10:11.
- Chalmers DJ (1995) *The conscious mind. Search of a Fundamental Theory* (Oxford Univ Press, New York).
- Nisbett RE, Wilson TD (1977) Telling more than we can know—Verbal reports on mental processes. *Psychol Rev* 84:231–259.
- Keyers C, Kaas JH, Gazzola V (2010) Somatosensation in social perception. *Nat Rev Neurosci* 11:417–428.
- Niedenthal PM (2007) Embodying emotion. *Science* 316:1002–1005.
- Goldstone R (1994) An efficient method for obtaining similarity data. *Behav Res Methods Instrum Comput* 26:381–386.
- Suvilehto JT, Glerean E, Dunbar RIM, Hari R, Nummenmaa L (2015) Topography of social touching depends on emotional bonds between humans. *Proc Natl Acad Sci USA* 112:13811–13816, and erratum (2015) 112:E6718.
- Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD (2011) NeuroSynth: A new platform for large-scale automated synthesis of human functional neuroimaging data. *Frontiers of Neuroinformatics Conference Abstract: 4th INCF Congress of Neuroinformatics*, 10.3389/conf.fninf.2011.08.00058.
- Han L, Kashyap AL, Finin T, Mayfield J, Weese J (2013) UMBC_EBIQUITY-CORE: Semantic textual similarity systems. *Second Joint Conference on Lexical and Computational Semantics* (Assoc Comput Linguist, Atlanta).