Mapping the semantic organization of the English odor vocabulary using natural language data

Hörberg, Thomas[1,2], Larsson, Maria[1] & Olofsson, Jonas[1]
[1]Gösta Ekman Laboratory, Department of Psychology, Stockholm University
[2]Department of Linguistics, Stockholm University

Olfactory experiences are hard to verbalize, partly because most languages lack devoted odor vocabularies. Yet, there is a need for a standardized odor vocabulary, but no descriptive system for describing the full range of odor experiences has been agreed upon. Many studies of the English odor vocabulary have been based on perceptual data such as odor-descriptor ratings, thereby being limited to a small set of pre-selected descriptors. In the present study, we present a data-driven approach that automatically identifies odor descriptors in English, and then derive their semantic organization on the basis of their distributions in natural texts. Olfactory descriptors are automatically identified on the basis of their degree of olfactory association, and their semantic organization is derived with a distributional-semantic word embedding model. We identify and derive the semantic organization of the descriptors most frequently used to describe odors and flavors in English, both within and across source-based, abstract and evaluative descriptor categories. Our method is to a large extent able to capture semantic differences between descriptors related to aroma and flavor qualities, rather than e.g. functional or linguistic aspects, in that it primarily differentiates descriptors with respect to valence and edibility, and the semantic space derived from it is qualitatively similar to a space derived from perceptual data.

key words: odor semantics, odor categorization, distributional semantics, word embeddings, corpus linguistics

Introduction

The human sense of smell has often been assumed to be vestigial in comparison to other senses. People are poor at discriminating and identifying familiar odors (e.g., De Wijk & Cain, 1994; Engen, 1991). In particular, people find it hard to name odors (Herz & Engen, 1996; Jönsson & Stevenson, 2014), and olfaction has even been described as "the muted sense" (Olofsson & Gottfried, 2015). With some notable exceptions (e.g. Burenhult & Majid, 2011; Wnuk & Majid, 2014), most languages lack a vocabulary that primarily is used to denote odor qualities. Instead of using devoted terms for describing odors (similar to, e.g., color terms), speakers often use either *source-based* descriptions, identifying odors on the basis of its source (e.g., *citrusy*), *abstract* descriptions, referring to some abstract (e.g., *musty*) or metaphorically used cross-modal sensory property (e.g., *light*), or *evaluative* descriptions, describing hedonic evaluation of the odor (e.g., *pleasant*) (Burenhult & Majid, 2011; Croijmans & Majid, 2016; Majid et al., 2018; Majid & Burenhult, 2014). Olfaction has therefore been found to be one of the senses that is hardest to verbalize (e.g., in terms of poor linguistic codability, see Majid et al., 2018).

Yet, in many industries and businesses, there is a need for a standardized vocabulary for classifying the perceptual space of odors and facilitating objective communication about smells. Although various olfactory taxonomies have been proposed within specific domains (e.g. Burlingame et al., 2004; Fisher et al., 2018; Jellinek, 1997; Noble et al., 1987; Suffet & Rosenfeld, 2007), no descriptive system for describing a wider variety of odor qualities has been agreed upon (Kaeppler & Mueller, 2013).

In this work, we present a study on the semantic organization of the English odor vocabulary, using an approach that is based on natural language data. Based on the distribution of words in olfactory and gustatory contexts in a three-billion-word corpus of written English, we identify olfactory descriptors by quantifying the degree of olfactory-semantic content they convey. We then derive their semantic

organization using a distributional-semantic word embedding model. In this model, semantic distances between words are represented as vector distances in a multi-dimensional space. Crucially, our method provides a fully data-driven approach for identifying a large set of olfactory descriptors and their semantic organization that is based on the distribution of words in texts from the Internet. We further investigate the dimensions of our semantic space using measures of word valence, word edibility associations and word concreteness, and finally validate our method against perceptual data in the Dravnieks odor-term rating data set (Dravnieks, 1985).

Several studies have investigated the perceptual and semantic space of odors using quantitative and multivariate analyses of perceptual data such as the Dravnieks data set. Khan et al. (2007) used principal component analysis (PCA) to reduce the dimensionality of the Dravnieks data. They found the primary dimension (PC1) to reflect odor pleasantness, a finding that was confirmed in a series of perceptual rating tasks. Similarly, Zarzo (2008) found the primary dimension (PC1) of odor-descriptor ratings similar to the Dravnieks data to be related to pleasantness, and the second dimension (PC2) to reflect edibility. Koulakov et al. (2011) found that the positions of the odors in the Dravnieks data set in a three-dimensional space could be approximated by a 2-dimensional curved manifold (explaining 52% of the variance in the data) whose primary dimension (the elevation of the manifold) was associated with odor pleasantness. Castro et al. (2013) derived a 10-dimensional perceptual space from the Dravnieks data using non-negative matrix factorization, a dimensionality reduction technique that constrains the principal dimensions to be positive. They found each dimension to be sparsely occupied by a small set of odors, and each odor profile to primarily be characterized by a single dimension. They interpreted these dimensions as meta-descriptors (e.g., 'fragrant', 'woody', 'fruity', and 'sickening'), reflecting fundamental and distinct perceptual qualities of individual odors. Importantly, the identified dimensions corresponded either to "relative pleasantness" or to "cues of potential palatibility / nonpalatibility" (Castro et al., 2013: 13). More recently, Licon et al. (2018) investigated the semantic space of odors on the basis of PCA analysis of odor-descriptor ratings. Participants rated 105 odors with respect to whether and how well they were described by 24 predefined descriptors[1], and with regard to their perceptual dimensions in terms of pleasantness, edibility, familiarity and intensity. In line with previous studies, they found the primary dimension (PC1) in their semantic space to be most highly correlated with pleasantness ratings, and the secondary dimension (PC2) to be primarily correlated with edibility ratings. Together, these studies provide plenty of evidence for the idea that odors are semantically and perceptually differentiated primarily with respect to pleasantness, on the one hand, and edibility, on the other.

There is only a handful of recent studies that have investigated the semantic and perceptual space of odors on the basis of natural language data. Some studies have used pre-trained distributional-semantic word embedding models to predict the applicability of odor descriptors to odor molecules. Nozaki & Nakamoto (2018) trained a neural network model to predict odor descriptor categories from chemical properties of odorants. Descriptor categories were identified on the basis of hierarchical clustering of cosine distances between descriptor vectors, derived from a pre-trained distributional-semantic word embedding model. Their model was able to predict one of 20 descriptor clusters with a 53% accuracy. Gutiérrez et al. (2018) took advantage of the correspondence between a perceptual and a semantic odor space, the former derived from perceptual odor-descriptor ratings, and the latter from pre-trained semantic vectors, to predict perceptual ratings from a small set of 19 descriptors in the DREAM data set (Keller et al., 2017) to 131 descriptors in the Dravnieks data set. This was done both for an identical set of odor molecules available in both data sets (i.e., their "direct semantics" model), as well as for a novel set of molecules only available in the Dravnieks data set (i.e., their "imputed semantics" model), for which DREAM descriptor ratings were estimated from chemoinformatic features of the shared odor molecules. Thus, Gutiérrez et al. (2018) showed that perceptual ratings for novel chemicals can be estimated from semantic vector representations combined with molecular information. These studies indicate that, for a given set of odors, there will be a high correspondence

---

[1] The choice of descriptors was done on the basis of several odor atlases and the comprehensive review of Zarzo and Stanton (2009).

between their perceptual space derived from odor-descriptor ratings, and their semantic space derived from a distributional-semantic word embedding model. Their applicability is, however, limited to a pre-selected set of descriptors, either from which categories are identified, or for which odor-descriptor ratings are readily available.

Other studies have shown how natural language data can be used to quantify words and phrases with respect to their olfactory and gustatory association and specificity. Whereas some descriptors most frequently are used to express odor or flavor qualities (e.g., *smelly*), others are more general and often used in other contexts (e.g., *sharp*). Further, some descriptors are used specifically for a limited set of odors or flavors (e.g., *flowery*), whereas others apply to a wider range of qualities (e.g. *strong*). In order to identify the vocabulary associated with wine qualities in a corpus of wine reviews, Crojimans et al. (2019) employed Termhood analysis, a method that automatically identifies domain-specific vocabulary on the basis of the degree to which words or phrases are used in specific contexts in comparison to their general use. They identified 146 unique terms that were used to construct a text-based wine wheel. Iatropoulos et al. (2018) proposed a corpus-based measure of olfactory / gustatory association (similar to Termhood analysis), the *Olfactory Association Index* (OAI). OAI is the log2 probability that descriptor d occurs in a olfactory-related context as opposed to a non-olfactory context, i.e.,

$$OAI_d = \log_2 \frac{of_d}{tf_d}$$

where tf_d is the total frequency of d and of_d is the frequency of d in olfactory-related contexts. OAI was found to be strongly correlated ($\rho = .693$) with experimental ratings of olfactory association, collected by Lynott & Connell (2009), attesting to the viability of OAI as a measure of olfactory association. Iatropoulos et al. (2018) further devised a measure of olfactory / gustatory specificity, the *Olfactory Specificity Index* (OSI), defined as the negative log2-probability that descriptor d occurs in a olfactory-related context,

$$OSI_d = -\log_2 \frac{of_d}{\sum_{\forall d \in D} of_d}$$

where D is the set of all descriptors investigated. As such, OSI estimates the information contained in a descriptor about olfactory contexts. High OSI indicates a high information content in terms of the descriptor being used in a limited set of contexts, and a low OSI is indicative of low information content in that the descriptor is used across a broad range of contexts. OSI was found to be negatively correlated ($r = -.250$) with inter-individual inconsistency of odor descriptor use, quantified in terms of descriptor-to-odor rating variance across subjects. Thus, olfactory descriptors that score low on OSI are used more inconsistently across speakers, presumably because they are used for a broader range of odor percepts and therefore allow for a wider usage across individuals.

Although these methods provides a data-driven way to identify descriptors used to describe odor and flavor qualities, they are uninformative regarding how those descriptors are organized semantically. In the present work, we overcome this by first identifying sets of descriptors with the highest olfactory association, using OAI. We then derive the semantic organization of these descriptors on the basis of word vectors from a distributional-semantic word embedding model. Importantly, in order to capture the meanings of the descriptors as applied to odor and flavor qualities, we use word vectors from a language model that is specifically trained on olfactory and gustatory contexts.

Method

Corpus

We identified olfactory descriptors and trained our language model using the University of Maryland Baltimore County (UMBC) webbase corpus (Han et al., 2013). The UMBC corpus consists of English written texts of more than three billion words in 100 million web pages from over 50,000 websites. We used the part-of-speech (POS) tagged version of the corpus with POS tags in the Penn Treebank format (Marcus et al., 1994; Taylor et al., 2003). We used the POS tagging to disambiguate between homonyms from different word classes.

The corpus was pre-processed by converting all words to lower case, lemmatizing content words (using the NLTK python module, see Bird et al., 2009), and dividing texts into sentences. We also identified 116 two-word or three-word collocations (shown in the supplementary materials), consisting of phrases either used in the Dravnieks data set (e.g., "black pepper", "rotten fruit", "nail polish remover") or considered to be commonly used to describe flavors and odors (e.g., "bad breath", "essential oil", "sweet and sour"), that were merged.

From the preprocessed corpus, we extracted an olfactory sub-corpus of around 7.9 million words in 293 960 olfactory and gustatory contexts. This was done by identifying sentences that contained any of 105 olfactory- or gustatory-related key words (e.g., 'odor', 'fragrance', 'taste', 'savor', the full list of words is shown in the supplementary materials).

Descriptor identification

From the olfactory corpus, we extracted all nouns and adjectives that occurred at least 10 times within a ±4 word olfactory/gustatory context window of the olfactory/gustatory key words, resulting in the initial identification of 7010 words. OAI and OSI of these words was calculated on the basis of olfactory/gustatory frequency and total frequency, and their corresponding lemmas / concepts were extracted (e.g., using "CITRUS" as lemma for "citrusy" or "citrus-like"). Among the words, we identified a number of word categories that we considered as unsuitable descriptors that were removed. These included technical terms, proper names, foreign terms (mainly latin), parts of lexicalized compounds, abbreviations, and terms that were included due to errors in the corpus (e.g., wrong POS tag or repeated texts). A subset of 2122 words of the remaining words were then manually classified as source-based, abstract, or evaluative. To the extent possible, this was done in accordance with the guidelines in the supplementary materials of Majid et al. (2018). However, some descriptors, such as hyperonyms (e.g., "fruit") and basic category adjective derivatives (e.g., "citrusy"), could potentially be classified as either source-based or abstract. Here, we chose to classify hyperonym nouns (e.g., "fruit") as source-based, as these refer to a specific category. All basic category terms (e.g., "citrus" and "citrusy") were classified as source-based, as these refer either to a specific object *per se* or the quality of that specific object. Hyperonym adjective derivatives (e.g., "fruity"), on the other hand, were classified as abstract, as these refer to a categorical quality, rather than the quality of a specific object.

Measures of word valence, edibility and concreteness

In order to further investigate the dimensions along which the descriptors primarily are differentiated, we took advantage of freely available data sets containing ratings of words on different semantic dimensions. Given previous findings, we wanted to see whether and to what extent our identified dimensions corresponded to pleasantness / valence, on the one hand, and to edibility, on the other.

For word pleasantness / valence, we used the word valence ratings provided by Warriner et al. (2013). This data set consists of ratings of approximately 14 000 word lemmas, rated in terms of valence, arousal and pleasantness on 1-9 grade lickert scale. For valence, participants rated words with respect

to how *happy* or *unhappy* they made them feel. Participants were informed that low ratings corresponded to feelings of being *happy, pleased, satisfied, contented* or *hopeful*, and high ratings to feelings of being *unhappy, annoyed, unsatisfied, melancholic, despaired, or bored* (Warriner et al., 2013:1193). We extracted valence ratings for those words that overlapped with the lemmas of the descriptors used in our study. In total, 1 717 - around 81% - of our descriptors were assigned a valence rating.

As en estimate of word edibility, we used the two data sets containing ratings of the perceived strength of association between words and sensory modalities, provided by Lynott & Connell (2009, 2013). They had participants rate 400 nouns (Lynott & Connell, 2013) and 423 adjectives (Lynott & Connell, 2009) on a 1-5 grade lickert scale with respect to the extent to which they associated each word with the visual, auditory, tactile, olfactory and gustatory sensory modality. Again, olfactory and gustatory ratings for words that overlapped with the lemmas of our descriptors were extracted. However, this resulted in the assignment of ratings to only 272 of our descriptors, making up about 13%. Therefore, we only include these ratings in our analyses in cases where ratings are available for at least 25% of the included descriptors.

We were also interested in investigating whether any dimension corresponded to word *concreteness*, that is, the extent to which a word denotes a perceptual entity (Brysbaert et al., 2014). Importantly, concreteness is highly related to the distinction between source-based, abstract and evaluative descriptors, evaluative descriptors such as *pleasant* being abstract, and source-based descriptors such as *apple* being concrete. Concreteness is also related to the distinction between nouns and adjectives, the former mainly denoting perceptual entities, and the latter frequently referring to abstract properties. We used the concreteness ratings of Bryesbart et al. (2014). This data set consists of ratings of 37 058 words of different word classes, that were rated on 1-5 point lickert scale. Of these, 1931 corresponded to the lemmas of our descriptors, making up about 91% of all of our descriptors.

The correlations between the variables are shown in Table 1. As can be seen in the table, olfactory and gustatory ratings are highly correlated with each other, as well as with OAI. This suggests that OAI can serve as an indirect measure of edibility.

Table 1. Correlations between descriptor rating variables, based on data for 272 descriptors for which all ratings are available. P-values are corrected for multiple comparisons on the basis of the method proposed by Holm (1979). Shadings highlight the strongest correlation of each variable.

Variable	OAI		OSI		Valence		Gust.		Olf.		Concr.	
	r	p	r	p	r	p	r	p	r	p	r	p
OAI	-	-	-0.03	1.00	0.07	1.00	0.66	<.0001	0.58	<.0001	-0.02	1.00
OSI	-0.03	1.00	-	-	-0.10	1.00	-0.17	.160	-0.20	.040	0.09	1.00
Valence	0.07	1.00	-0.10	1.00	-	-	0.20	.040	0.10	1.00	0.07	1.00
Gust.	0.66	<.0001	-0.17	0.16	0.20	.040	-	-	0.65	<.0001	-0.04	1.00
Olf.	0.58	<.0001	-0.20	.036	0.10	1.00	0.65	<.0001	-	-	0.07	1.00
Concr.	-0.02	1.00	0.09	1.00	0.07	1.00	-0.04	1.00	0.07	1.00	-	-

Language model

The sentences in the olfactory corpus was used to train a distributional-semantic word embedding model. In this model, semantic distances between words are represented as vector distances in a multi-dimensional space. The semantic vectors are the hidden layers of a 3-level feedforward neural network trained to predict a target word on the basis of its context. Words that occur in similar contexts have similar meanings and therefore also similar vector representations. The model was trained with the continuous-bag-of-words (CBOW) algorithm of word2vec, using the Python gensim package

(Řehůřek & Sojka, 2010). We used a context-window of ±4 words, a hidden layer consisting of 200 units, a minimum word frequency count of 10, and 30 training iterations.

Data pre-processing

We extracted word vectors from our language model that corresponded to subsets of our identified descriptors (see below), on the one hand, and to the set of source-based descriptors used in the Dravnieks data set, on the other. Each word vector was standardized and from the standardized word vectors, descriptor distance matrices were derived. The distance D between descriptor i and j was calculated as

$$D_{i,j} = 0.5 \times (1 - \rho_{i,j})$$

As such, the distance D is the pearson (ρ) correlation between word vectors, converted to the 0-1 range, with 0 reflecting semantic identity (i.e, corresponding to $\rho = 1$) and 1 indicating semantic opposition (i.e., corresponding to $\rho = -1$).

A distance matrix was also derived on the basis of the standardized Dravnieks rating vectors of the selected Dravnieks descriptors (see below), using the same method.

Statistical analysis

In order to identify descriptor clusters and to further investigate the primary dimensions along which the descriptors are differentiated, the distance matrices were analyzed with hierarchical clustering and PCA. As clustering method, we used the Agglomerative Nesting (AGNES) algorithm (Kaufman & Rousseeuw, 1990), as implemented in the R package 'Cluster' (Maechler et al., 2019). AGNES is a bottom-up clustering method that successively merges clusters on the basis of their distance, quantified on the basis of Ward's minimal variance method (Ward, 1963). PCA analysis was conducted with the prcomp() function, implemented in the base R stats package ({R Core Development Team}, 2018). In order to provide stronger evidence for our interpretations of the principal components, we also investigate how well the components correlate with OAI and OSI, as well as with independent measures of word valence, concreteness and gustatory/olfactory association.

Results

Top descriptors combined

Initially, we investigated the semantic organization of the top source-based, abstract, and evaluative descriptors combined. To this end, we selected the top 50 source-based, the top 50 abstract, and the top 50 evaluative descriptors in terms of OAI score with at least 20 occurrences in the olfactory corpus[2]. Due to their ambiguous nature, basic category adjectives (e.g., "citrusy") were excluded in these analyses. For the selected 150 descriptors, we derived a distance matrix as described above.

Hierarchical cluster analysis

First, we analyzed the distance matrix with hierarchical clustering. The clustering is illustrated in Figure 1. The main partitioning differentiates between descriptors that mainly are used for unpleasant, strong or overpowering experiences, and descriptors that are used for pleasant or flavorful experiences. On the basis of the clustering, we further distinguished between five main clusters which contain semantically similar descriptors. It should be stressed that this partitioning primarily was done on the basis of our intuitions. In addition to our clustering, we also calculated the Dunn clustering Index for

---

[2] Many of the descriptors with less than 20 occurrences in the olfactory corpus were highly uncommon and limited to highly specific contexts, e.g., *mignonette*. Here, we focus on descriptors that are used in everyday language.
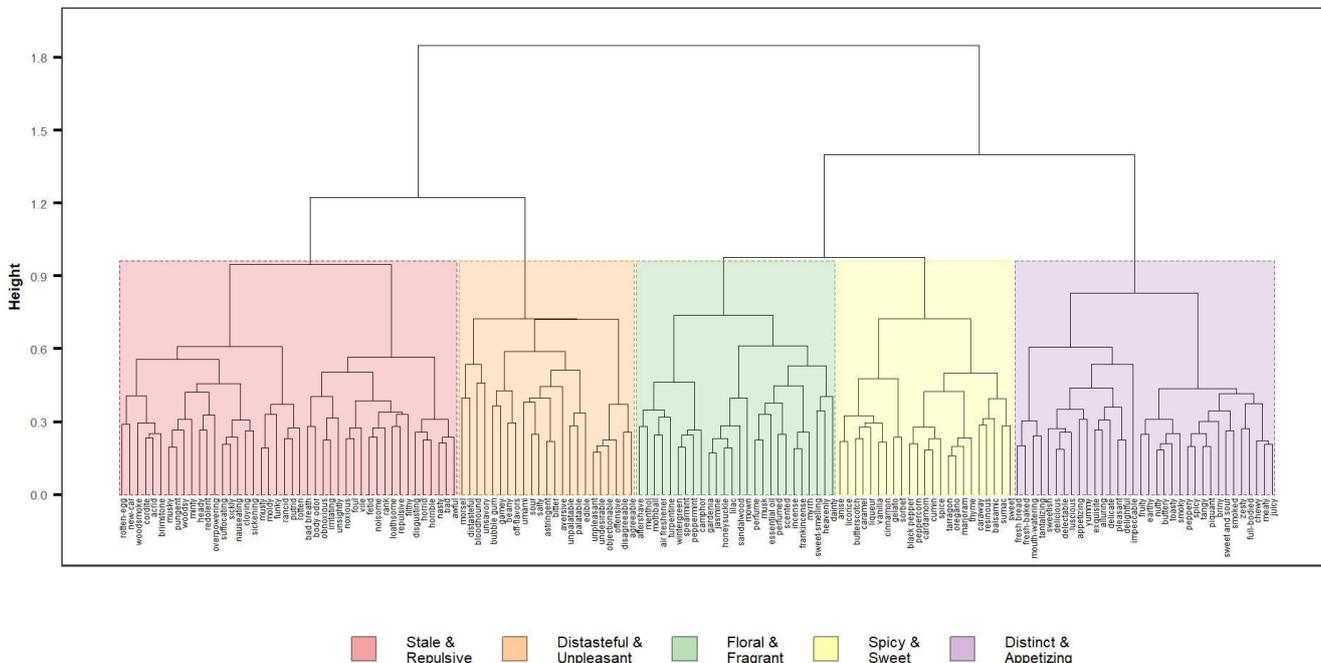
Figure 1. Hierarchical clustering of distance matrix of the top 50 source-based, the top 50 abstract, and the top 50 evaluative descriptors.

between 4 to 15 cluster partitions. The Dunn Index is the ratio of the smallest distance between observations not in the same cluster to the largest intra-cluster distance (Dunn, 1974). A higher value corresponds to a more optimal clustering solution in terms of keeping across-cluster items separated. This analysis found 10 cluster partitions to be optimal (Dunn's Index: 0.17). Thus, our five clusters each contain two sub-clusters which together make up the optimal clustering solution.

The first of our five clusters, referred to as *Stale & Repulsive* in Figure 1, primarily contains abstract and evaluative descriptors often used to describe highly unpleasant odors such as *stale, repulsive, pungent* and *disgusting*. It also contains a few source-based descriptors that may be used for unpleasant odors such as *rotten-egg, woodsmoke,* and *cordite*. The second cluster, *Distasteful & Unpleasant*, primarily contains abstract and evaluative descriptors related to (un)palatability and unpleasant flavors (e.g., *palatable, distasteful, offensive*), and descriptors denoting basic flavors, which, importantly, often are used to describe negative flavor experiences (e.g., *bitter, sour, salty*). The third cluster, *Floral & Fragrant,* contains both source-based and abstract descriptors denoting minty (e.g., *menthol, spearmint, peppermint*), floral (e.g., *jasmine, lilac, floral*) and fragrant / aromatic (e.g., *perfumed, musk, incense*) odors. The fourth cluster, *Spicy & Sweet*, contains source-based and abstract descriptors denoting sweet (e.g., *caramel, vanilla, licorice*), spicy (e.g. *black pepper, peppercorn, sumac)* and herbal (e.g., *tarragon, marjoram*) aromas. The fifth and final cluster, *Distinct & Appetizing,* contains abstract and evaluative descriptors denoting highly positive aroma and flavor experiences (e.g., *delicious, exquisite, delightful*), on the one hand, and specific and in many cases positive aroma and flavor experiences (e.g., *fruity, spicy, juicy*), on the other. Importantly, this descriptor clustering seems to a large extent to reflect semantic differences between descriptors that involve aroma and flavor qualities, rather than, being based on, for instance, functional (e.g., clustering food products together) or categorical (e.g., clustering all source-based descriptors together) criteria. This indicates that our word embedding model is able to capture and distinguish between the aroma and flavor qualities that the descriptors denote.

PCA analysis

In the PCA analysis of the distance matrix, the first four PCs together explain more than 70% of the variance in the data (32.09%, 17.23%, 12.37% and 8.79%, respectively), and we therefore restrict our analysis to those. In Figure 2, the normalized loadings and contributions of the top 30 and bottom 30 descriptors of each component are illustrated. Table 2 shows the correlations between the four PCs,
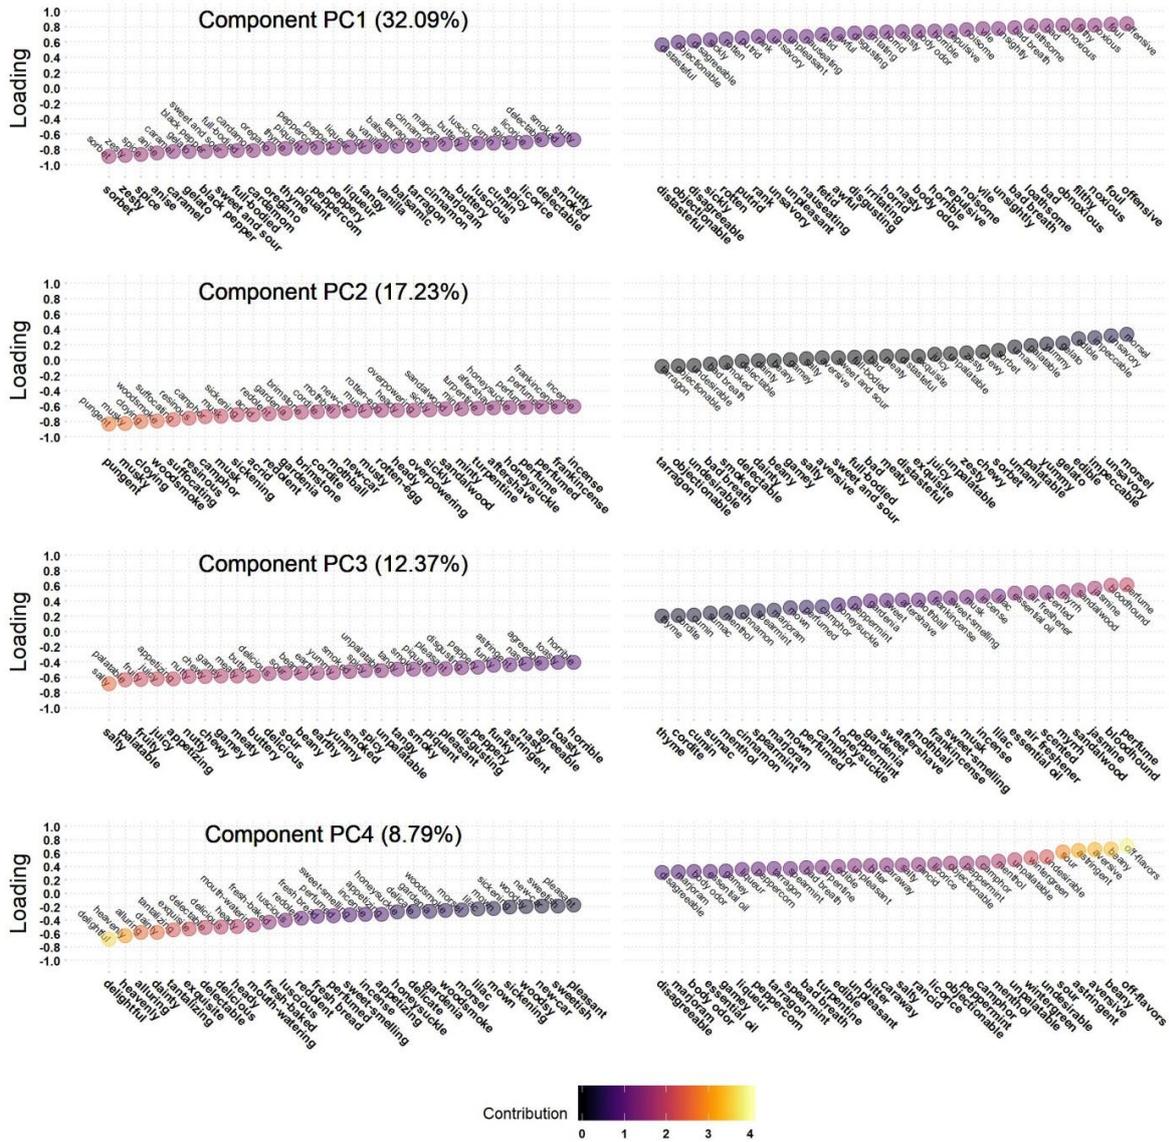
Figure 2. Normalized component loadings and contributions of the top 30 and bottom 30 descriptors of each component of the PCA model of the top 150 descriptors.

OAI and OSI, as well as the valence, edibility and concreteness ratings, and Figure 3 illustrates these correlations. Taken together, our results show that PC1 primarily is related to Valence / Pleasantness, Edibility / Gustation, but also to a somewhat less extent to Concreteness. More specifically, whereas descriptors with high scores on PC1 tend to primarily be related to pleasant odors and flavors (e.g.,

Table 2. Correlations between the first four principal components of the PCA model of the top 150 descriptors, on the one hand, and OAI / OSI and the semantic variables, on the other. P-values are corrected for multiple comparisons on the basis of the method proposed by Holm (1979). Shadings highlight the strongest correlation of each variable.

| Variable	N	PC1 (32.09%)		PC2 (17.23%)		PC3 (12.37%)		PC4 (8.79%)	
		r	*p*	r	*p*	r	*p*	r	*P*
OAI	150	0.32	.002	0.43	<.0001	-0.03	1.00	-0.21	.150
OSI	150	0.19	.222	-0.10	1.00	-0.28	.009	-0.06	1.00
Valence	101	0.73	<.0001	-0.04	1.00	-0.20	.530	0.29	.058
Gustation	41	0.63	<.0001	-0.32	.487	0.41	.111	-0.09	1.00
Olfaction	41	-0.01	1.00	0.27	.808	-0.17	1.00	0.19	1.00
Concreteness	126	0.51	<.0001	0.18	.530	-0.54	<.0001	-0.31	.009
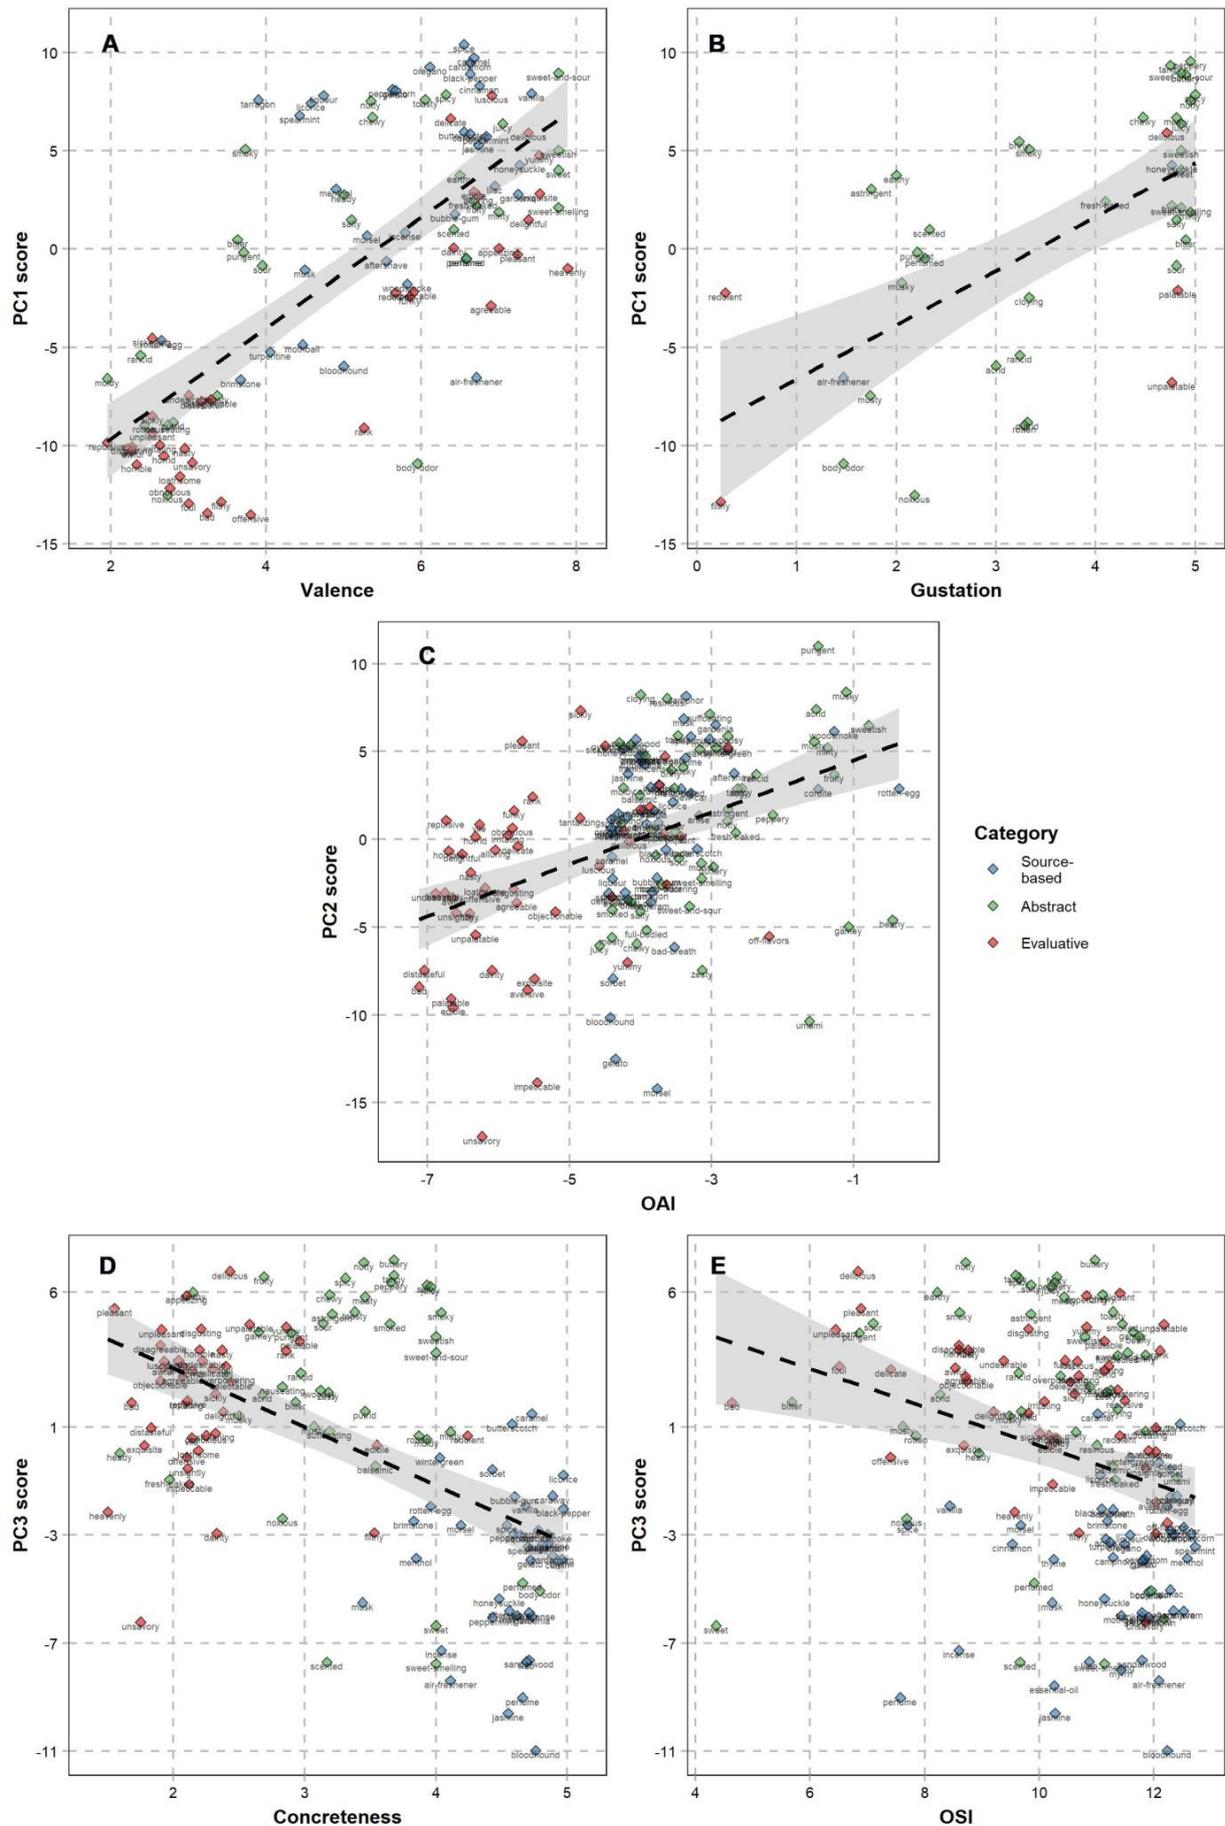
Figure 3. Correlations between descriptor PCs in the top 150 descriptor PCA and variables, with descriptors differentiated on the basis of descriptor type. A. PC1 and Valence. B. PC1 and Gustation. C. PC2 and OAI. D. PC3 and Concreteness. E. PC3 and OSI.

sorbet, caramel, luscious), low scoring descriptors mainly express unpleasant odors qualities or odor sources (e.g., *offensive, bad breath, putrid*). High scoring descriptors also tend to be either source-based or abstract and refer to edibles or flavor qualities, whereas low scoring descriptors most frequently are abstract or evaluative, referring to unpleasant odor qualities more generally. In line with these observations, PC1 is highly correlated to both evaluative ratings and gustatory association ratings (see Table 2 and Figure 3), but also to concreteness.

These findings are in line with the results of previous studies investigating the semantic organization of olfactory descriptors on the basis of rating data (e.g. Khan et al., 2007; Zarzo, 2008) in that they show that odor descriptors primarily are differentiated in terms of pleasantness and edibility. This finding thus provide initial support for the viability of mapping the organization of the odor lexicon on the basis of word embeddings. PC2, on the other hand, is clearly not related to valence. Descriptors expressing both pleasant (e.g., *perfumed, aftershave*) as well as unpleasant (e.g., *pungent, sickening*) odor experiences have high negative loadings on PC2. Instead, as shown by the fairly strong correlation between PC2 and OAI, PC2 is related to degree of olfactory association. Descriptors that are strongly associated with olfaction have high negative loadings on PC2. For PC3, it is mainly source-based and concrete descriptors that have high positive loadings (e.g., *perfume, jasmine, mothball*), and abstract and evaluative descriptors that have high negative loadings (e.g., *salty, palatable, disgusting*). Thus, as further shown by the fact that PC3 is highly correlated to concreteness ratings but none of the other variables, PC3 is primarily related to descriptor concreteness. Consequently, whereas mainly source-based descriptors score high on PC3, abstract and evaluative descriptors score low. PC4, finally, is also negatively correlated with concreteness ratings. However, with all three types of descriptors evenly distributed across PC4, it is harder to make any interpretation regarding PC4.

In Figure 4, the semantic distribution of the descriptors are further illustrated in PC1, PC2 and PC3 score plots. Figure 4 also illustrates the main correlations between PCs and the variables.

Top descriptors of each category separated

We then further explored the distribution of descriptors within descriptor categories, by selecting descriptors with the highest OAI scores and at least 20 occurrences from each category separately. We selected the top 150 source-based and the top 150 abstract descriptors, this time with basic category adjectives included. Since there are a lot fewer evaluative descriptors in our data set, all evaluative descriptors with at least 20 occurrences were included, resulting in the selection of 85 descriptors. For the selected descriptor sets, we again derived three separate distance matrices as described above. In the following, we describe the analyses of each descriptor category in turn, starting with source-based descriptors, abstract descriptors and finally evaluative descriptors.

Source-based descriptors

Hierarchical clustering. Again, we first analyzed the source-based descriptor distance matrix with hierarchical clustering. This clustering is illustrated in Figure 5. Interestingly, in this analysis, the main partitioning mainly seems to differentiate between descriptors referring to unedible sources, on the one hand, and descriptors referring to edibles, on the other. The first main partitioning contains descriptors referring to unpleasant, stale and chemical sources, on the one hand, and flowery and fragrant sources, on the other. The second partitioning contains descriptors referring to fruits and sweets, herbs and spices, and other edibles. Again, as based on our intuitions, we further distinguished between five main clusters with semantically similar descriptors. Dunn's index (>0.12) indicated an optimal clustering solution of 11 clusters or more. However, our five-cluster-partitioning received the second highest score of 0.1197. The first source-based cluster, *Stale & Chemical*, primarily contains descriptors referring to highly unpleasant odor sources and chemicals (e.g., *rotten egg, bad breath, piss* and *amonia*). The second cluster, *Flowery & Fragrant*, on the other hand contains descriptors referring to flowery, pleasant and fragrant odor sources (e.g., *deodorant, perfume, lavender* and *rose*).
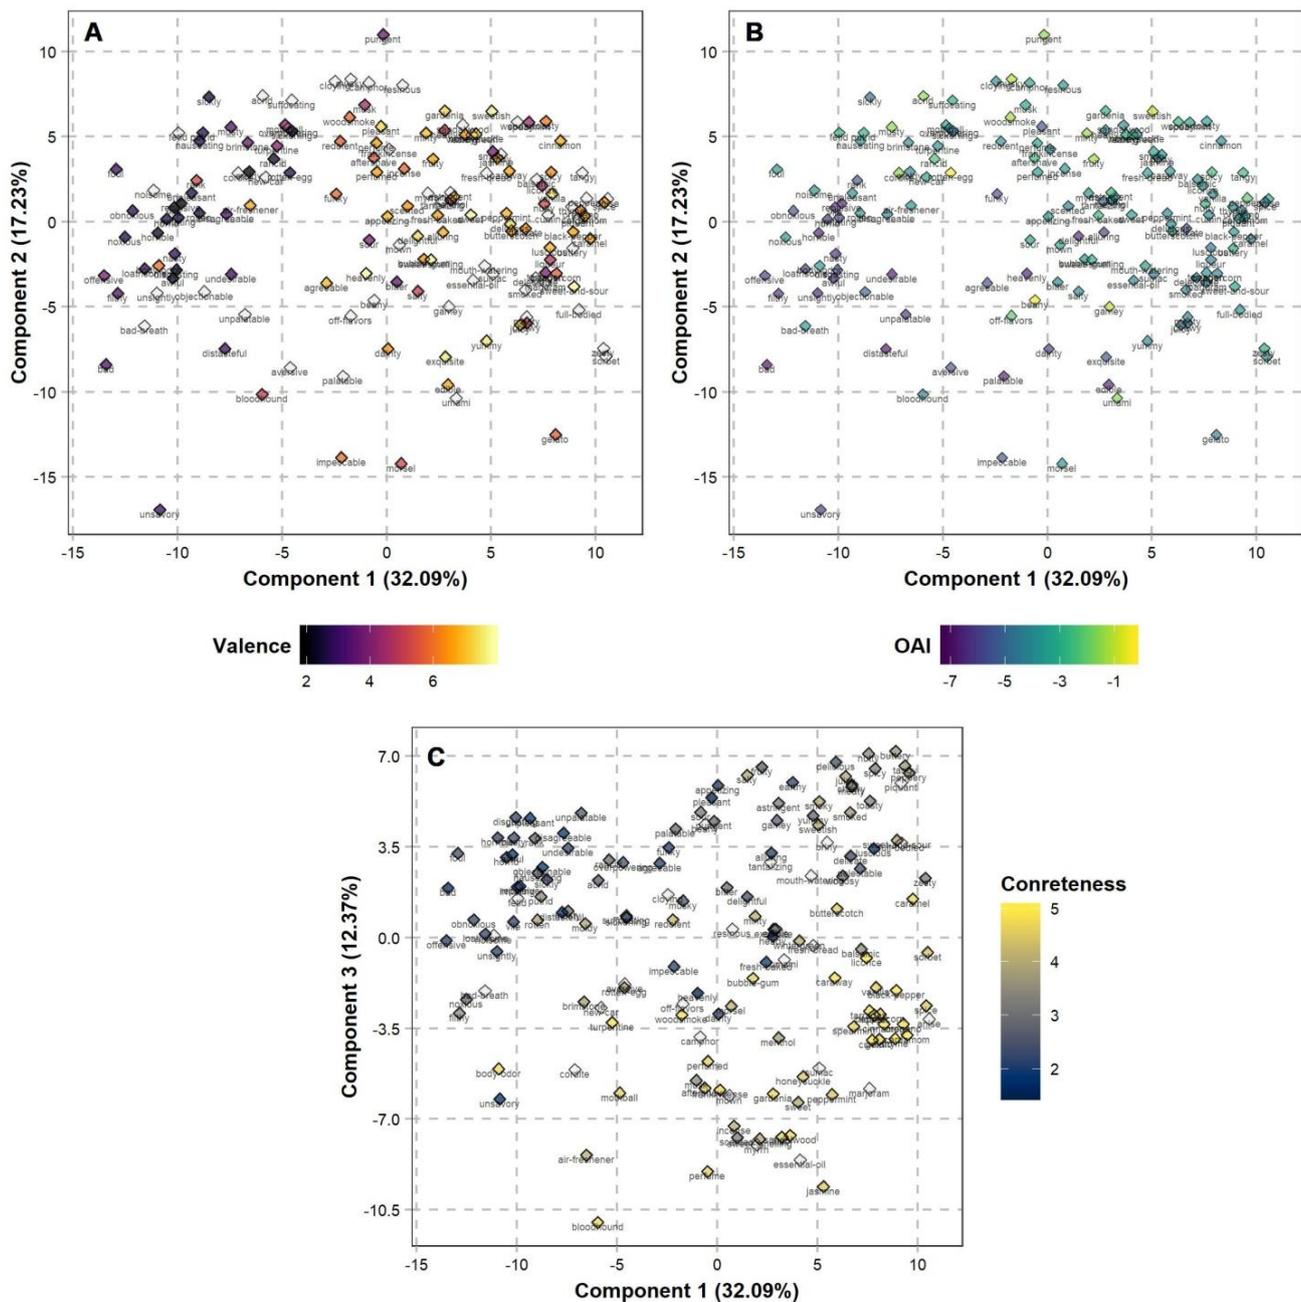
Figure 4. PCA score plots of PC scores of the top 150 descriptor PCA model, also showing variable scores. A. Descriptors for which no variable score was available are shown in light gray. A. PC1 and PC2 score plot together with Valence rating scores. B. PC1 and PC2 score plot with OAI. C. PC1 and PC3 with concreteness scores and descriptor type.

What these two clusters have in common is that they both contain descriptors referring to unedible odor sources. The third cluster, *Fruity & Sweet*, contains descriptors primarily referring to fruits (e.g., *strawberry*, *mango*), sweeteners (e.g., *caramel, honey, syrup*) and some other flavorful ingredients (e.g., *anise, mint, peanut butter*). The fourth cluster, *Herbal & Spicy*, primarily contains descriptors referring to herbs and spices (e.g., *black pepper, salt, oregano, thyme*). The fifth and final cluster, *Edible*, contains descriptors referring to heavy, meaty and savory ingredients (e.g., *truffle, cheese, venison*), dishes (e.g., *dumpling, risotto, kebab*), desserts (e.g., *cheesecake, sorbet, gelato*) or food-related hyperonyms (e.g., *morsel, treat, appetizer*).
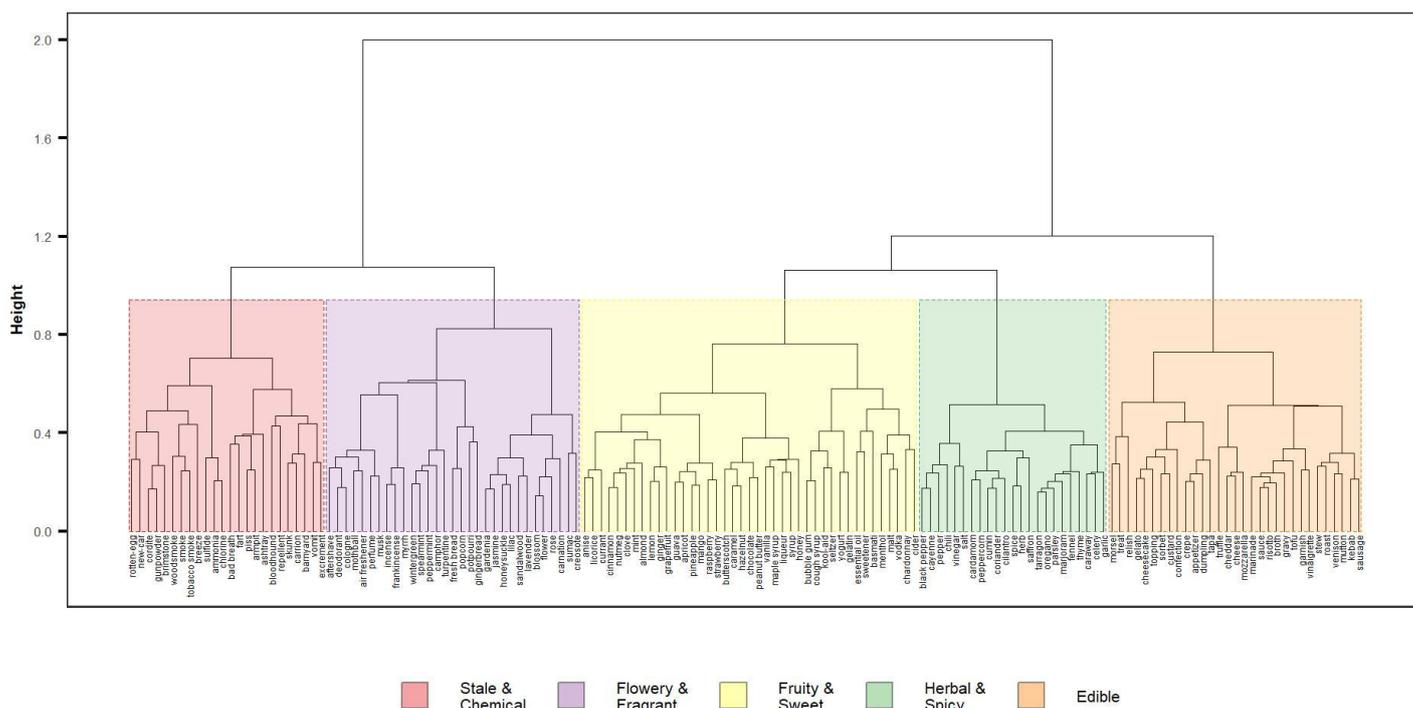
Figure 5. Hierarchical clustering of distance matrix of the top 150 source-based descriptors.

PCA analysis. We chose to investigate the first three PCs, that together explain 64.8% of the variance in the data; 38.78%, 18.52%, and 7.5% respectively. Figure 6 shows the normalized loadings and contributions of the top 30 and bottom 30 descriptors of each component, Table 3 shows the correlations between the three PCs and the variables (i.e., OAI, OSI, valence, edibility and concreteness), and Figure 7 illustrates these correlations.

Descriptors with high scores on PC1 tend to refer to odor sources with pleasant smells (*chocolate, apricot, strawberry*), and descriptors with low scores to odor sources with unpleasant smells (e.g., *excrement*, *armpit*, *ashtray*). In line with this, PC1 is highly positively correlated with valence ratings, and to some extent negatively correlated with OAI (see Figure 7).

We note, however, that it primarily is source-based descriptors referring to edibles that score the highest on PC1, and that many of the highly pleasant but unedible descriptors have lower scores. For example, flower descriptors (e.g., *lilac, jasmine, flower*) have *negative* PC1 scores. Thus, considering that valence and edibility are correlated (see Table 1), it could be that PC1 either conflates valence and edibility, or, that it primarily is related to edibility, and that it correlates with valence ratings due to the correlation between valence and edibility. Importantly though, it is clear that PC1 captures the same qualitative aspects of odor descriptors that also are found in ratings-based accounts (e.g., Kahn et al. 2007). However, PC2 seems to be more straightforwardly interpreted as reflecting edibility. Food-

Table 3. Correlations between the first three principal components of the PCA model of the top 150 source-based descriptors, on the one hand, and OAI / OSI and the semantic variables, on the other. P-values are corrected for multiple comparisons on the basis of the method proposed by Holm (1979). Shadings highlight the strongest correlation of each variable. The number of ratings available on each variable is also shown.

| Variable	n	PC1 (38.78%)		PC2 (18.52%)		PC3 (7.5%)	
		r	*p*	r	*p*	r	*p*
OAI	150	-0.25	.024	0.41	<.0001	-0.03	1.00
OSI	150	-0.02	1.00	-0.12	0.930	-0.07	1.00
Valence	121	0.45	<.0001	0.09	1.00	0.07	1.00
Concreteness	134	0.15	.654	0.16	.654	-0.05	1.00

12
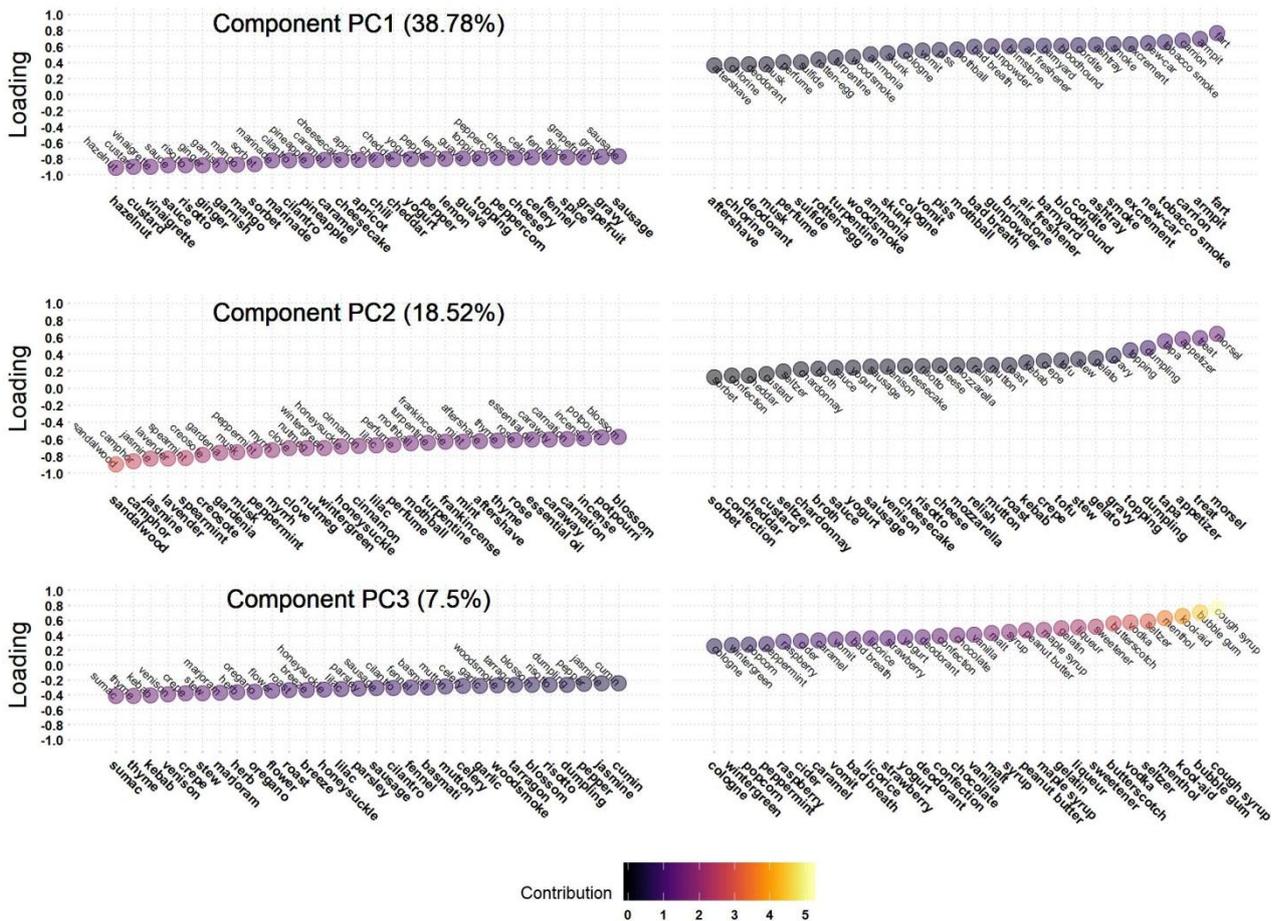
Figure 6. Normalized component loadings and contributions of the top 30 and bottom 30 descriptors of each component of the PCA model of the top 150 source-based descriptors.

related descriptors (e.g., *treat*, *appetizer*, *dumpling*) score high on PC2, and descriptors referring to flowery, fresh and fragrant odor sources (e.g., *jasmine*, *musk, perfume*) score low. Descriptors referring to unpleasantly smelling odor sources, generally have low PC2 loadings with scores close to zero. Further, PC2 is unrelated to pleasantness/valence, but positively correlated with OAI, which, in turn is highly positively correlated with edibility (see Table 1). PC2 therefore more clearly seems to be related to edibility. It is harder to interpret the meaning of PC3, since it does not correlate with any of the variables. However, descriptors referring to fresh (*bubble gum*, *kool-aid*, *menthol*) and sweet (e.g., *sweetner, maple syrup, strawberry*) sources score high on PC3, and descriptors referring to heavy and salty sources (e.g., *kebab, venison, crepe*), indicating that PC3 to some extent is related to a light-heavy continuum. In sum, whereas PC1 primarily is related to pleasantness /valence, and to some extent also to edibility, PC2 is primarily related to edibility and independent of pleasantness.

Abstract descriptors

Hierarchical clustering. The results of the hierarchical clustering of the abstract distance matrix is shown in Figure 8. Again, the clustering primarily differentiates between descriptors that mainly are used for unpleasant experiences, on the one hand, and descriptors used for pleasant experiences, on the other. We further separated these two main partitioning into five main clusters with semantically similar descriptors. In the first partitioning, we identified two distinct clusters. The first cluster, referred to as *Distinct & Toxic* in Figure 8, consists of descriptors referring to basic flavors (e.g., *gamey, oily, bitter, sour*), often used for negative experiences, descriptors referring to toxic experiences (*noxious, poisonous, toxic*), as well as perception-related descriptors (*detectible, perceptible, noticeable*). The second cluster, referred to as *Stale* in Figure 8, contains descriptors that refers to directly unpleasant and stale experiences (e.g., *rancid, pungent, stale, putrid*). In the second
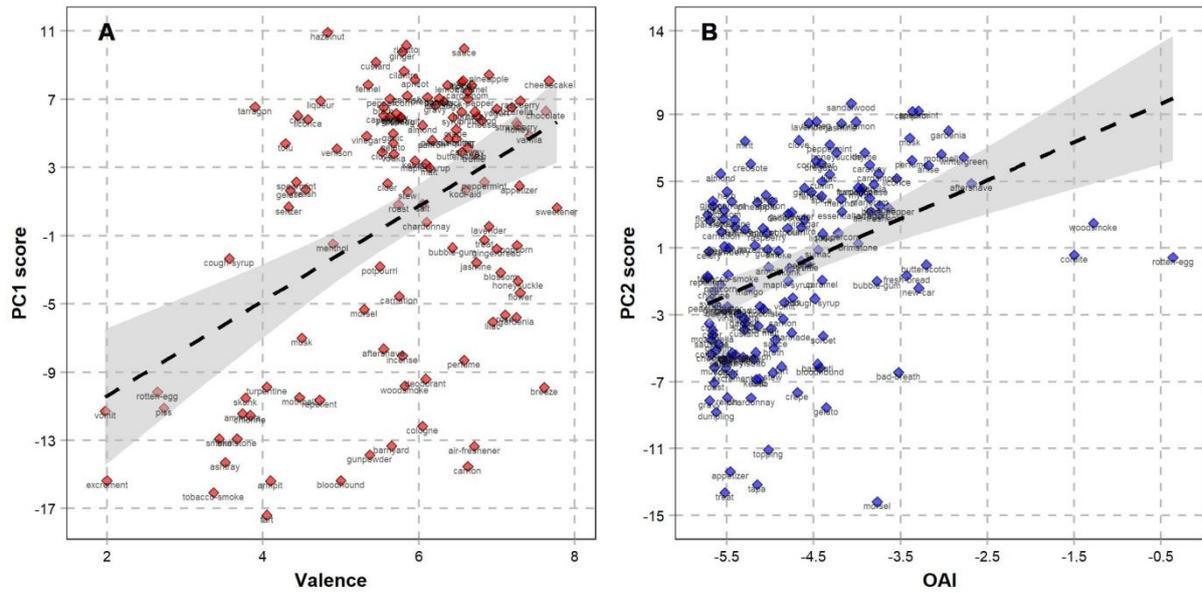
Figure 7. Correlations between the PCs in the top 150 source-based descriptor PCA and variables. A. PC1 and Valence. B. PC1 and OAI.

partitioning with pleasant descriptors we identified three main clusters. The first, labeled *Sweet, Floral & Fragrant,* contains descriptors referring to sweet, minty and fresh experiences (e.g., *sugary, minty, refreshing*) as well as perfumey, fragrant, strong and sensual experiences (e.g. *bridal, perfumed, unmistakable, sensuous*). The second cluster, *Edible*, contains descriptors referring to food-related and culinary experiences (e.g. *nutty, cheesy, crispy* and *zesty*). The fifth and final cluster, *Refined*, contains descriptors referring to complex and cultural experiences often considered to be of a higher standard (e.g., *aristocratic*, *aesthetic* and *eclectic*). For the abstract descriptors, the optimal cluster solution as based on Dunn's index consists of 10 clusters (Dunn's Index: 0.15).
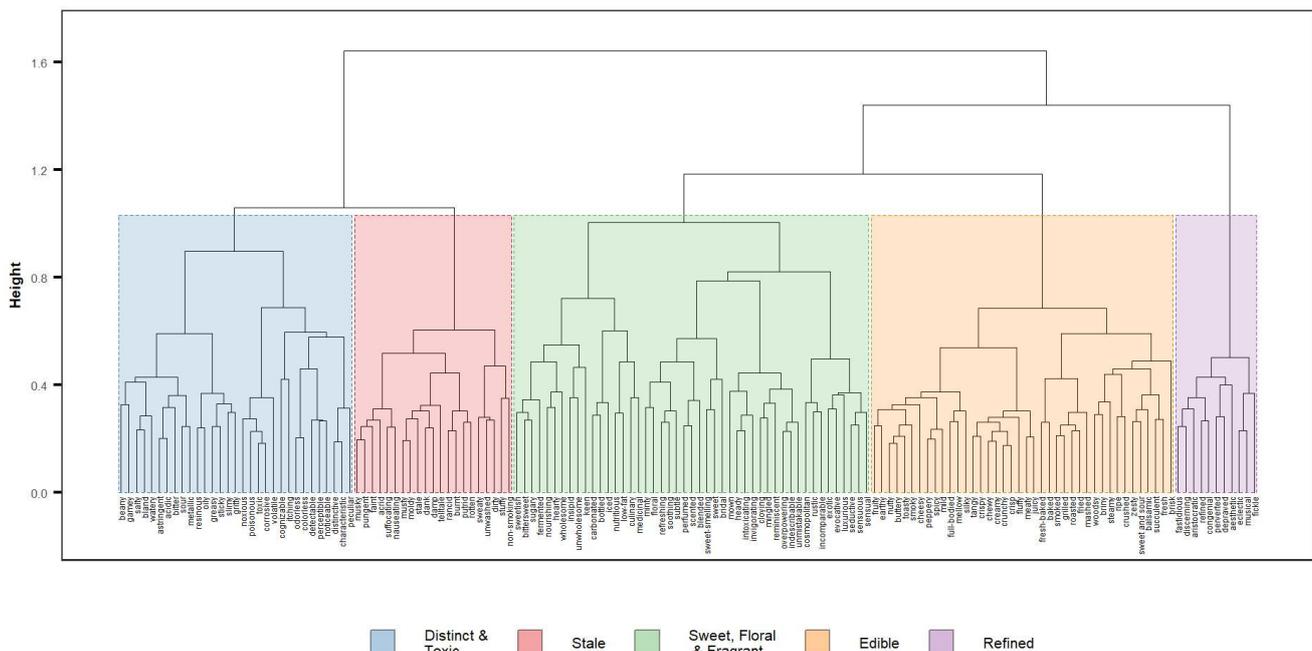
Figure 8. Hierarchical clustering of distance matrix of the top 150 abstract descriptors.

14

*PCA analysis.* For the PCA analysis, we further investigated the first four PCs, explaining more than 67% of the variance (30.42%, 20.57%, 9.41%, and 7.07%, respectively). The normalized loadings and contributions of the top 30 and bottom 30 abstract descriptors of each component are shown in Figure 9. Table 4 shows the PC and variable correlations, and Figure 10 illustrates these correlations.

On PC1, it primarily is fairly concrete descriptors referring to food-related and culinary experiences (e.g., *cheesy, nutty, peppery*) that have high positive loadings and low scores, and a few descriptors expressing complex and cultural experiences (e.g., *aesthetic, musical, congenial*) that have negative loadings and high scores. PC1 is also negatively correlated with OAI and to some extent to concreteness. Thus, given the high correlation between OAI and gustatory ratings, PC1 seems to primarily be related to (un)edibility, with descriptors expressing fairly concrete, food-related experiences scoring low on PC1.
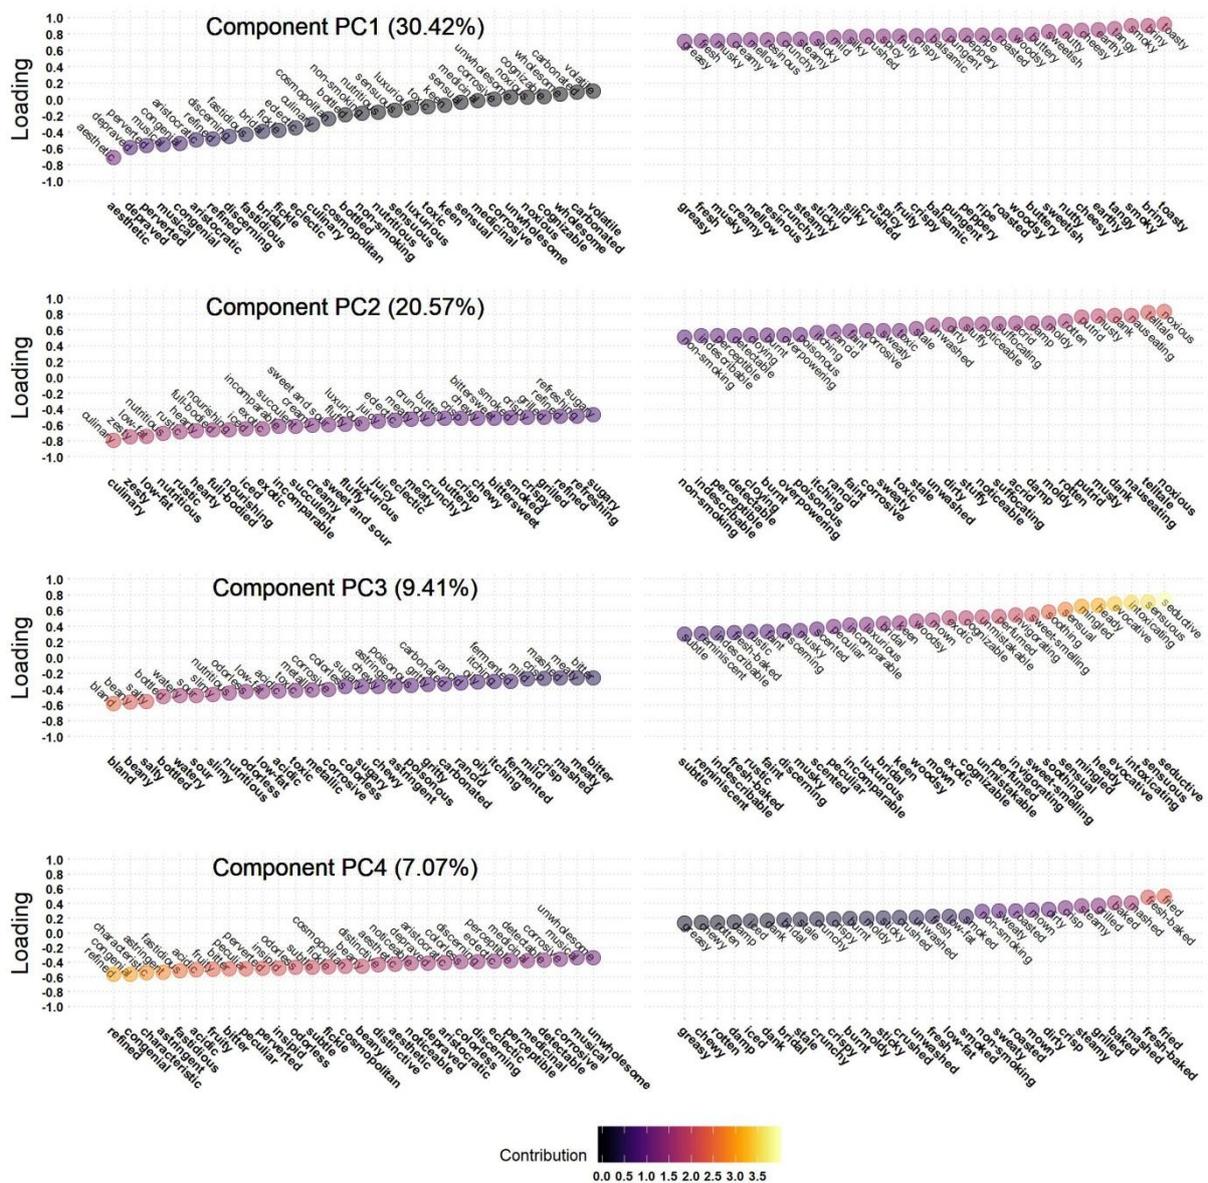
Figure 9. Normalized component loadings and contributions of the top 30 and bottom 30 descriptors of each component of the PCA model of the top 150 source-based descriptors.

Table 4. Correlations between the first three principal components of the PCA model of the top 150 abstract descriptors, on the one hand, and OAI / OSI and the semantic variables, on the other. P-values are corrected for multiple comparisons on the basis of the method proposed by Holm (1979). Shadings highlight the strongest correlation of each variable. The number of ratings available on each variable is also shown.

Variable	n	PC1 (30.42%)		PC2 (20.57%)		PC3 (9.41%)		PC4 (7.07%)	
		r	p	r	p	r	p	r	p
OAI	150	-0.48	<.0001	-0.11	1.00	0.01	1.00	0.06	1.00
OSI	150	0.06	1.00	0.14	1.00	-0.08	1.00	-0.13	1.00
Valence	103	0.01	1.00	0.67	<.0001	-0.42	<.0001	0.04	1.00
Gustation	62	-0.33	.152	0.40	.022	0.23	1.00	0.19	1.00
Olfaction	62	0.07	1.00	-0.29	.344	-0.42	.011	-0.04	1.00
Concreteness	136	-0.34	<.0001	-0.00	1.00	0.40	<.0001	-0.46	<.0001

For PC2, on the other hand, descriptors expressing unpleasant experiences have high positive loadings and low scores, and descriptors expressing sweet, refreshing, culinary and food-related and pleasant experiences have high negative loadings and high scores. PC2 also has a strong positive correlation with Valence, and is positively correlated with gustatory ratings. Thus, it seems clear that PC2 primarily is related to valence, but also to some extent to edibility, probably because abstract odor and flavor experiences that are pleasant also often tend to be culinary.

Descriptors expressing sensual, sweet, fragrant, and pleasantly complex experiences, often with an abstract meaning (e.g., *sensuous, intoxicating, invigorating*) have high positive loadings and low scores on PC3, and descriptors referring to basic flavors and food-related experiences, often interpreted as having a negative connotation (e.g., *salty*, *beany*, *bland*), unpleasant experiences, as well as more concrete properties (e.g., *bottled, slimy, carbonated*), have negative loadings and high scorese. PC3 is further negatively correlated with Valence and olfactory ratings, but positively correlated with Concreteness. Thus, PC3 seems to be a quite complex dimension differentiating between descriptors expressing highly pleasant odor and aroma experiences in abstract ways, and descriptors with more concrete meanings, describing neutral and sometimes negative, every-day odor and flavor experiences and food- and drink-related properties. For PC4, finally, rather concrete descriptors referring either to food-related experiences or other concrete properties have high loadings and low scores, and more abstract descriptors with more complex and metaphorical uses have low loadings and high scores. PC4 is also negatively correlated with Concreteness, and thus, similarly to PC3, seems to be related to concreteness.

Evaluative descriptors

Hierarchical clustering. Figure 11 shows the results of the hierarchical clustering of the evaluative distance matrix. Once again, it is clear that the clustering differentiates between descriptors that mainly referring to unpleasant experiences, on the one hand, and descriptors denoting pleasant experiences, on the other. Based on our intuitions, we further divided these two partitionings into four main clusters with semantically similar descriptors. This four-cluster partitioning received the second highest score on Dunn's index (.139), being very close to the score of the optimal clustering solution of 15 clusters (Dunn's index = .147). The first cluster, named *Unpleasant & Distasteful* in Figure 11, contains descriptors referring to highly unpleasant odor and flavor experiences (e.g., *foul, sickening, disgusting*). The second cluster, *Acceptable,* contains a few descriptors denoting unpleasant experiences (e.g., *unpleasant*, *offensive*, *aversive*), but also a few rather neutral terms (e.g., *acceptable*, *agreeable*, *edible*). In the third cluster, *Delicious*, there are descriptors expressing positive experiences (e.g., *pleasant, favorite, alluring*), many of which are flavor-related (e.g., *piquant, mouth-watering, appetizing*). The fourth and final cluster, *Magnificent*, contains descriptors expressing highly pleasant experiences that in many cases imply a high standard or top quality (e.g., *marvelous, magnificent, heavenly*).
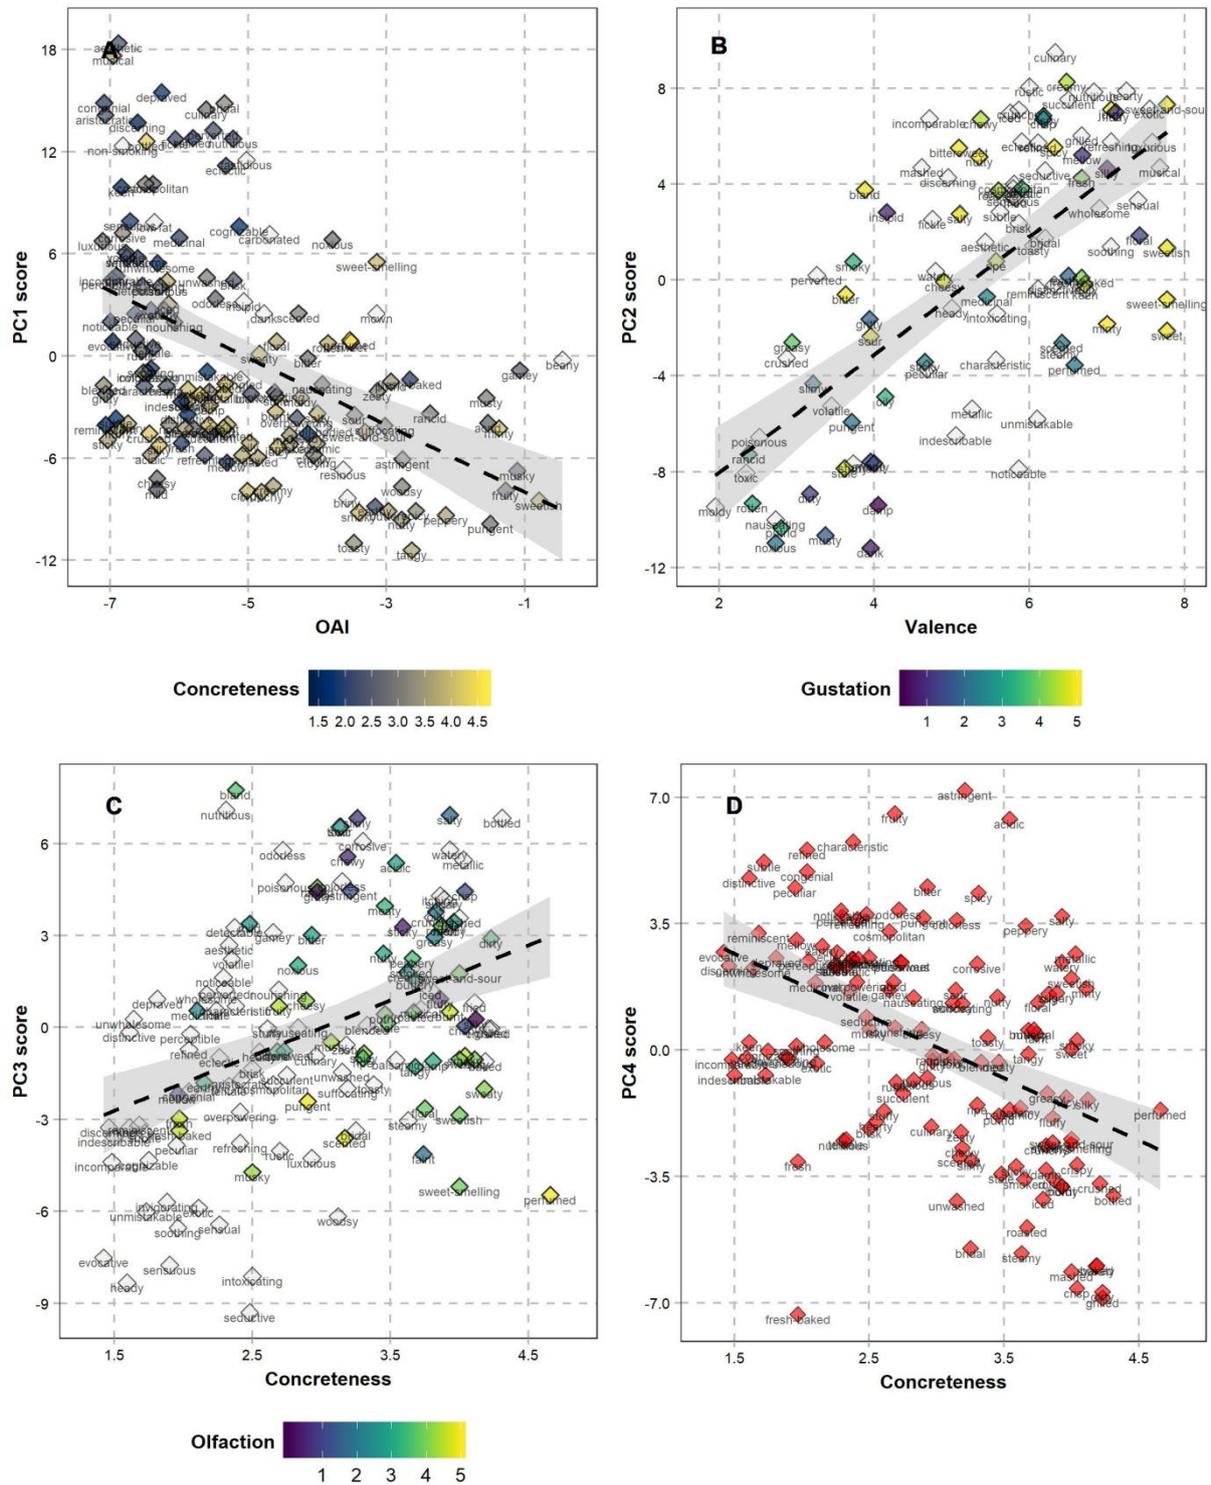
Figure 10. Correlations between the PCs in the top 150 abstract descriptor PCA model and variables, also showing secondary variable ratings. Descriptors for which no variable score was available are shown in light gray. A. PC1-OAI correlation together with Concreteness ratings. B. PC2-Valence correlation with Gustation ratings. C. PC3-Concreteness correlation with Olfaction ratings. D. PC4-Concreteness correlation.

PCA analysis. For the PCA analysis, we further investigated the first three PCs, explaining more than 64% of the variance (41.30%, 12.13%, and 10.59%, respectively). The normalized loadings and contributions of the top 30 and bottom 30 abstract descriptors of each component are shown in Figure 12, Table 5 shows the PC and variable correlations, and Figure 13 illustrates these correlations.
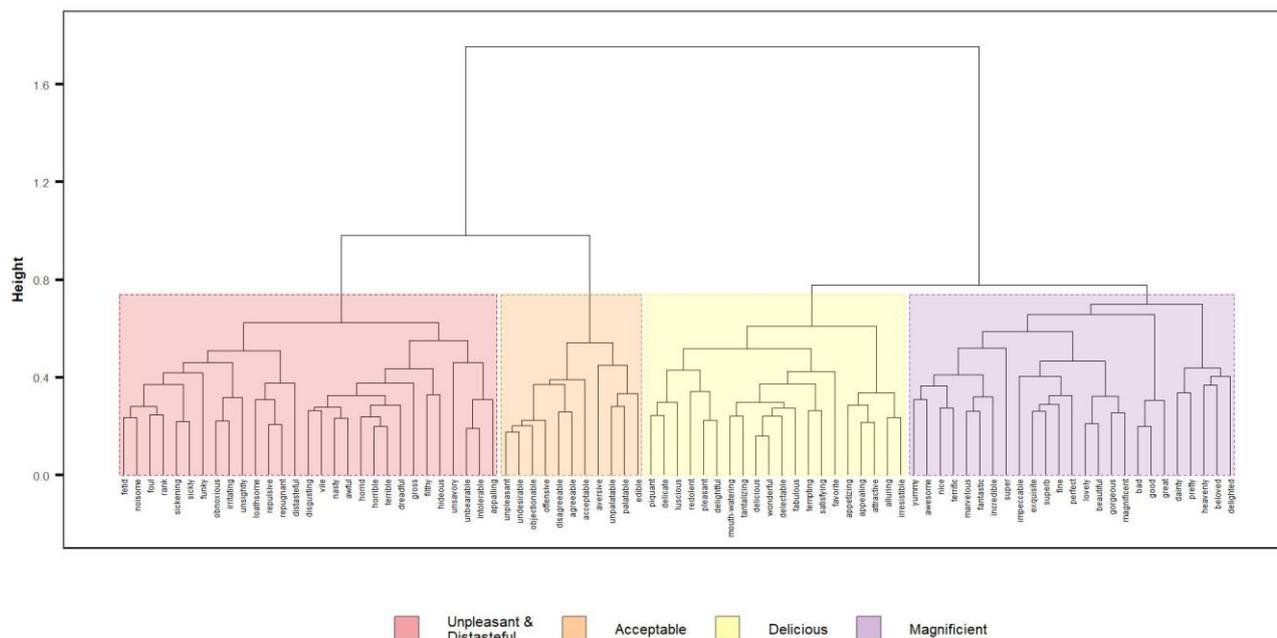
***Figure 11.*** Hierarchical clustering of distance matrix of the top evaluative descriptors.

PC1, explaining more than 40% of the variance of the data, is clearly related to pleasantness. Whereas descriptors expressing unpleasant experiences have high PC1 loadings and low PC1 scores, descriptors referring to high pleasantness and high quality experiences have low loadings and high scorese. Further, PC1 is highly positively correlated to Valence ratings. On PC2, on the other hand, descriptors expressing both pleasant and unpleasant experiences have high loadings and low scores. PC2 descriptor loadings appear to be fairly unsystematic, although many descriptors referring to pleasant food-related experiences (e.g., *delicious, appetizing, mouth-watering*) tend to have high loadings, and *unsavory* having a particularly low loading (see Figure 12). PC2 is also negatively correlated with OAI. Since OAI in turn is positively correlated with gustatory ratings, it could be that PC2 is related to a pleasantness cline with respect to food-related experiences. Descriptors with low PC2 scores and high PC2 loadings refer to pleasant food-related experiences (e.g., *delicious, appetizing, mouth-watering*). For PC3, finally, descriptors expressing pleasant as well as unpleasant experiences have high positive loadings without any apparent systematicity, and PC3 is uncorrelated with all of the variables. Thus, it is unclear what PC3 is related to.

Comparison between a ratings-based and word embedding-based descriptor space

As a final step, we compared a descriptor space derived from our word embedding model to a descriptor space derived from the perceptual data of the Dravnieks descriptor-odor ratings, using 145 source-based descriptors in the Dravnieks data (see supplementary materials for a list of all descriptors). For the selected descriptors, we extracted semantic vectors from the word embedding model, on the one hand, and descriptor-odor rating vectors from the Dravnieks data, on the other. On the basis of these (standardized) set of vectors, we calculated one ratings-based and one word embedding-based distance matrix. These distance matrices were then analyzed with hierarchical clustering and PCA. In the following, we compare the fit between these two analyzes, starting with the results of the hierarchical clustering.

Hierarchical clustering

The results of the hierarchical clustering is shown in Figure 14. The figure compares the ratings-based descriptor clustering (top panel) to the clustering based on the semantic vectors (bottom panel). The ratings-based clustering primarily differentiates descriptors referring to highly unpleasant odor sources (e.g., *garlic*, *feaces*, *cat urine*) from all other descriptors. However, this cluster also contains a couple

18
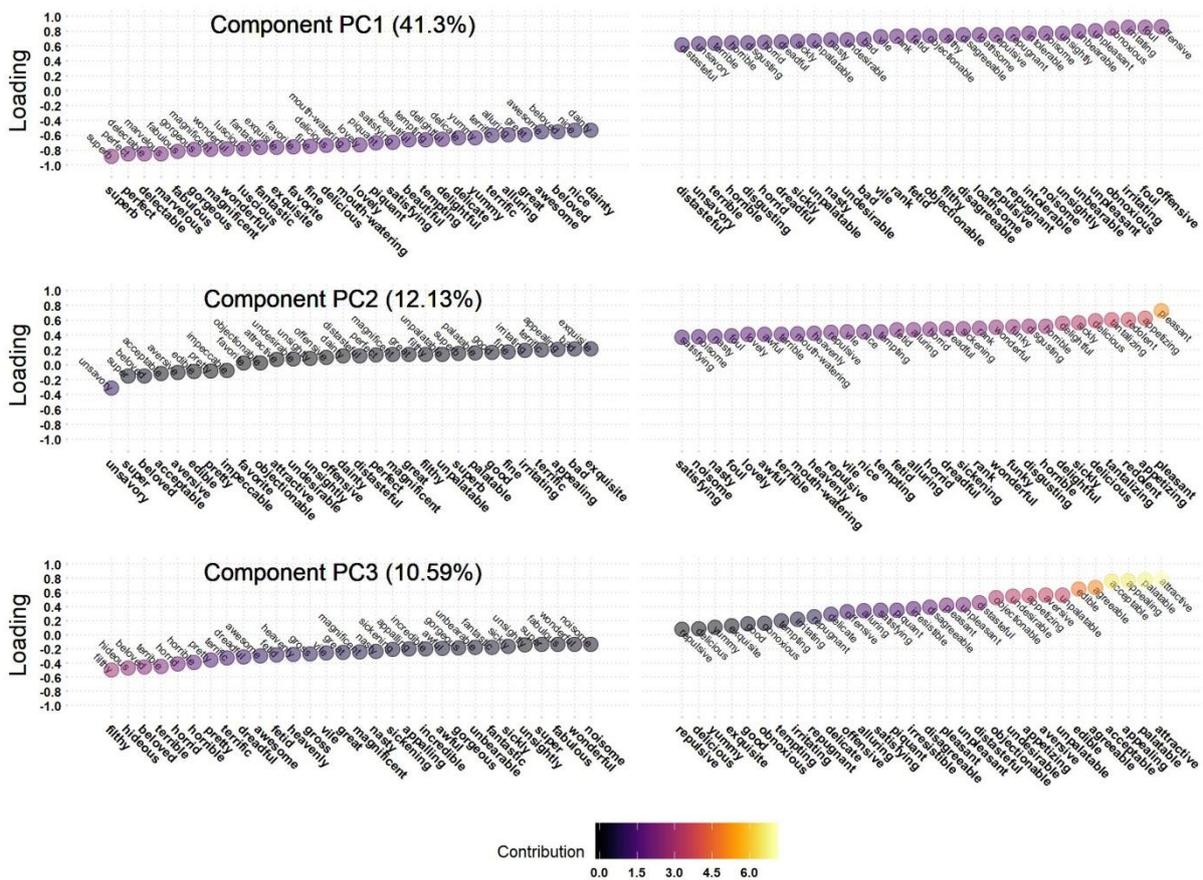
Figure 12. Normalized component coordinates and contributions of the top 30 and bottom 30 descriptors of each component of the PCA model of the top 85 evaluative descriptors.

of descriptors referring to edibles which not necessarily have an unpleasant smell (*soup*, *meat*, *fried chicken*), but nevertheless have fairly heavy and musty aromas. The word embedding clustering, on the other hand, primarily differentiates descriptors denoting edibles (e.g., s*trawberry, chocolate, potato*) from descriptors referring to unedibles (e.g., *rose*, *smoke*, *solvent*). These results indicate that whereas a semantic space of source-based descriptors that is derived from odor-descriptor ratings primarily is organized along the pleasantness dimension (in line with previous findings), a semantic space derived from semantic vectors of a word embedding model is first and foremost differentiated on the basis of edibility. With that being said, qualitatively, there is also several correspondences between the clusterings. Based on our intuitions, we further divided both clusterings into five main clusters containing semantically similar descriptors. In the ratings-based clustering, the first cluster, named *Fruity & Flowery* in Figure 14, contains descriptors referring to fruity (*peach, cherry*), flowery (*rose, violet*), as well as pleasant and fragrant (*cologne, perfume*) odor sources. The second cluster, *Aromatic & Spicy*, contains descriptors referring to sweet (*caramel, syrup*), aromatic (*chocolate,*

Table 5. Correlations between the first three principal components of the PCA model of the evaluative descriptors, on the one hand, and OAI / OSI and the semantic variables, on the other. P-values are corrected for multiple comparisons on the basis of the method proposed by Holm (1979). Shadings highlight the strongest correlation of each variable. The number of ratings available on each variable is also shown.

| Variable	n	PC1 (41.30%)		PC2 (12.13%)		PC3 (10.59%)	
		r	*p*	r	*p*	r	*p*
OAI	85	-0.26	.168	-0.38	.003	-0.21	.430
OSI	85	-0.06	1.00	0.22	.408	0.13	1.00
Valence	73	0.91	<.0001	0.15	1.00	-0.14	1.00
Concreteness	79	-0.08	1.00	0.01	1.00	-0.01	1.00
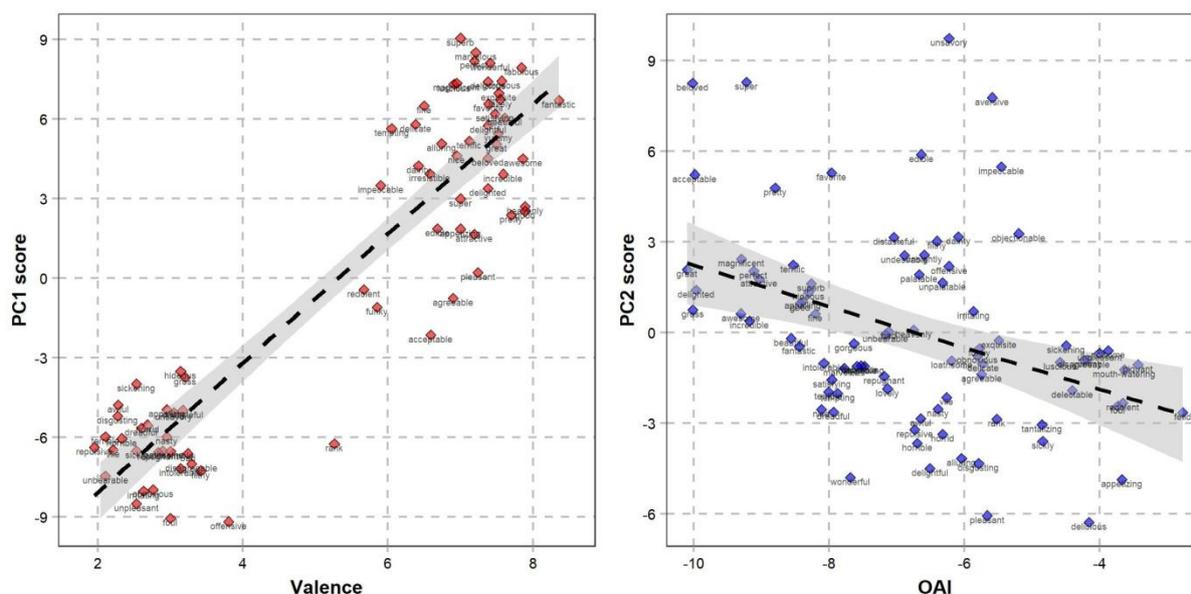
Figure 13. Correlations between the PCs in evaluative descriptor PCA model and variables. A. PC1-Valence correlation. B. PC2-OAI correlation.

coffee) and spicy (*cinnamon, anise*) sources. In comparison, the first cluster of the word embedding clustering (*Fruity & Aromatic* in Figure 14) contains descriptors referring to fruity, sweet and aromatic sources, and the second cluster, *Edible & Heavy,* contains descriptors denoting spicy sources as well as other (often heavy and savory) edibles (e.g., *bean, meat, cheese*). Thus, there are obvious correspondences between the two first clusters of each clustering solution. The third cluster of the ratings-based clustering, *Green & Heavy*, contains descriptors referring to (both edible and unedible) sources with heavy, powdery, and smokey odor qualities (e.g., *chalk, soot, cognac*), as well as descriptors referring to grassy, green and herbal sources (e.g., *herb, weed, dill*). In comparison, the third cluster of the word embedding clustering, *Flowery & Green*, contains descriptors denoting flowery, pleasant and fragrant odor sources, thus corresponding to the initial cluster of the ratings-based clustering. However, this cluster also contains descriptors referring to green and herbal sources, thus showing a partial overlap with the third cluster of the ratings-based clustering. There is also, finally, a quite clear correspondence between the two final clusters of the two clustering solutions. In the ratings-based clustering, the fourth cluster, *Chemical*, contains descriptors referring to mainly chemical and alcoholic odor sources (e.g., *ether, disinfectant, solvent*), and the fifth cluster, *Stale*, contains unpleasant (e.g., *garlic*, *feaces*, *cat urine*) and acidic (e.g., *acid, vinegar*) odor sources, as well as a few heavy and musty edibles (*soup*, *meat*, *fried chicken*). In the word embedding clustering, it is, conversely, the descriptors of the fourth cluster (*Stale*) that primarily refer to highly unpleasant sources (but also a few chemicals), whereas the descriptors of the fifth cluster (*Chemical*) mainly denote chemical, alcoholic and acidic sources.

In order to quantify the overlap between the two clustering solutions, we calculated the cophenetic correlation (Sokal & Rohlf, 1962), which, in our case, measures the correlation between the cophenetic distance matrices of our clusterings. We also tested the significance of the obtained correlation coefficient on the basis of a permutation test. This was done by randomly permuting the descriptor labels of the word embedding clustering 10 000 times, and calculating the cophenetic correlation between each of the permutated word embedding clusterings, on the one hand, and the ratings-based clustering, on the other. The distribution of the correlations from each permutation then serves as the probability distribution against which the observed correlation is tested; it approximates the probability distribution for the correlation that can be expected under the null hypothesis that the two clustering solutions are unrelated. The observed cophenetic correlation between the dendrograms is 0.268. Although this correlation is rather small, it is far larger than the correlation that is expected under the null hypothesis of the dendrograms being unrelated (range: -0.019, 0.068). Considering the

20
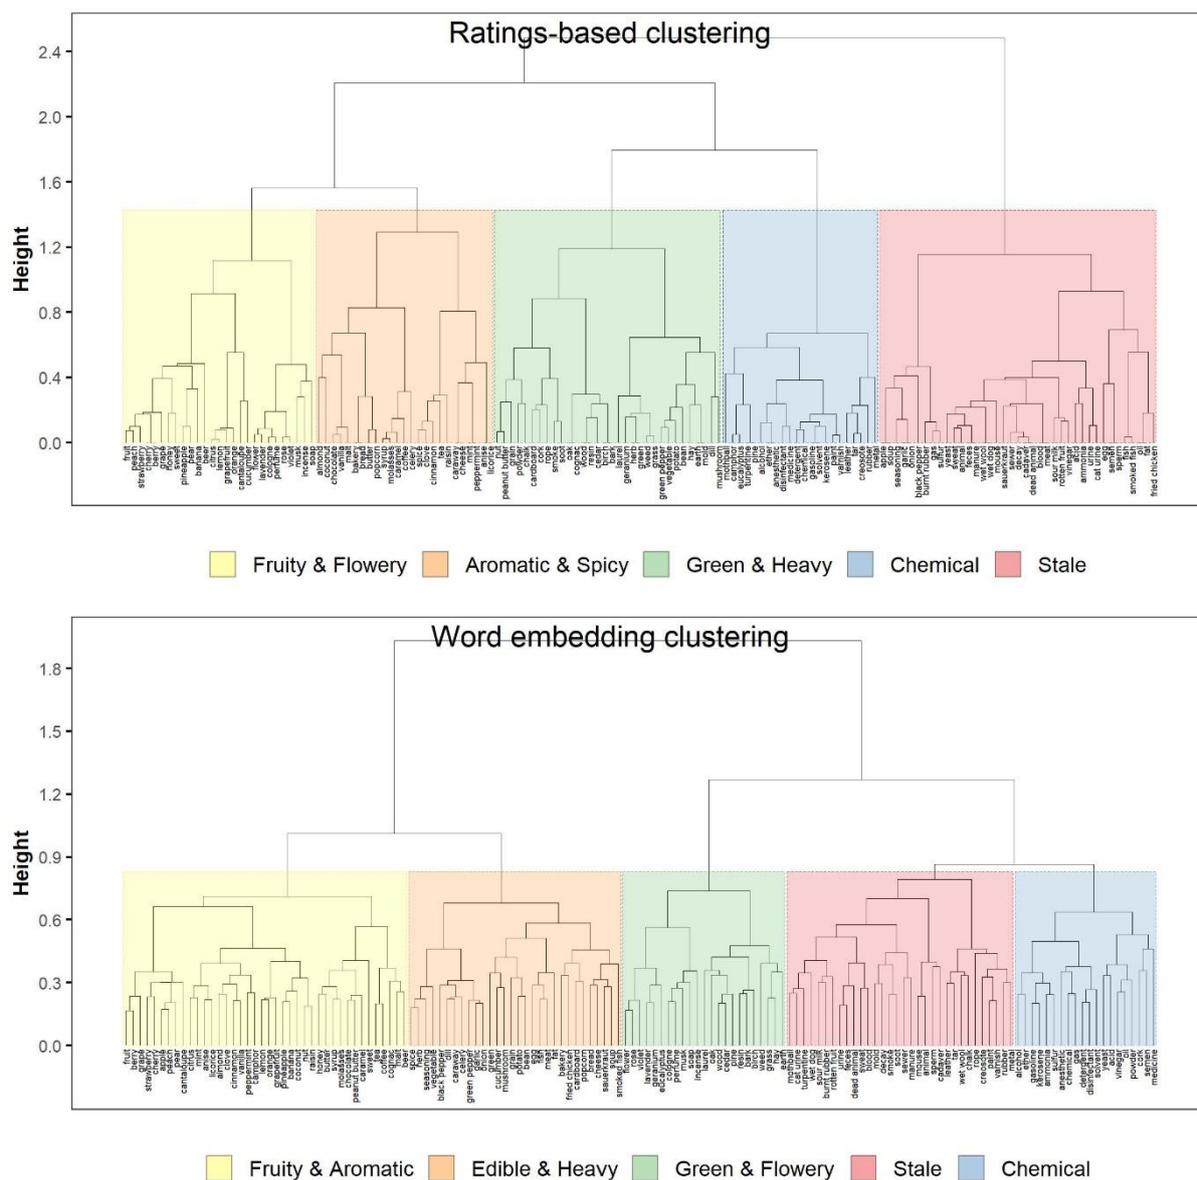
Figure 14. Hierarchical clustering of of the 146 source-based descriptors in the Dravnieks data set. Top panel: clustering performed on the ratings-based distance matrix. Bottom panel: clustering performed on the distance matrix based on word embeddings.

differences between the original datasets that the distance matrices are based upon (a collection of natural texts, on the one hand, and odor-descriptor ratings, on the other), these results are highly promising.

PCA analysis

For the results of the PCA analyses of the two distance matrices, we only investigate the first two principal components, as earlier work has shown the primary dimension of the Dravnieks data to reflect pleasantness, and the second to be related to edibility (e.g. Castro et al., 2013). We were particularly interested in whether this also is the case for the PCA space derived from the semantic vectors of our word embedding model, and therefore investigated the correlations between PC1 and PC2 of the word embedding distance matrix, on the one hand, and PC1 and PC2 of the ratings data, on the other.

Figure 15 illustrates PC1 and PC2 scores of the Dravnieks descriptors obtained in the PCA analysis of the ratings-based distance matrix. The figure also shows the clusterings identified by the hierarchical

clustering. Visual inspection of the figure largely confirm the findings of earlier studies (Khan et al., 2007; Zarzo, 2008). The primary dimension PC1, explaining 33.49% of the variance in the data, is primarily related to pleasantness. Whereas descriptors referring to pleasant odor experiences score low on PC1, descriptors expressing unpleasant experiences score high. The secondary dimension, PC2, explaining 19.45% of the variance in the data, seems to primarily be related to edibility; descriptors referring to edibles score high, and descriptors referring to non-edibles score low. Indeed, descriptors referring to chemical and alcoholic odor sources, that is, in many cases toxic substances, stand out with particularly low PC2 scores.

The correlation between PC1 and PC2 of the ratings-based PCA analysis, and PC1 and PC2 of the word embeddings-based PCA analysis is shown in Table 6, and further illustrated in Figure 16. Figure 17, finally, illustrates PC1 and PC2 scores of the word embedding PCA analysis, together with PC1 and PC2 scores of the ratings-based PCA analysis. As shown in Table 5, PC1 of the word embedding model is highly positively correlated with PC1 of the ratings-based model. However, it also has a moderate negative correlation with PC2 of the ratings-based model. As shown in Figure 16 and 17, descriptors referring to pleasant odor and flavor sources (e.g., *pineapple, fruit, sweet*) tend to have low scores on PC1 of the word embedding PCA model, and descriptors referring to unpleasant sources (e.g., *sewer, cadaver, dead animal*) tend to have high scores. However, it is also the case that descriptors referring to edibles tend to score low, and descriptors denoting unedibles tend to score high. Taken together, these findings indicate that while PC1 of the word embedding model primarily
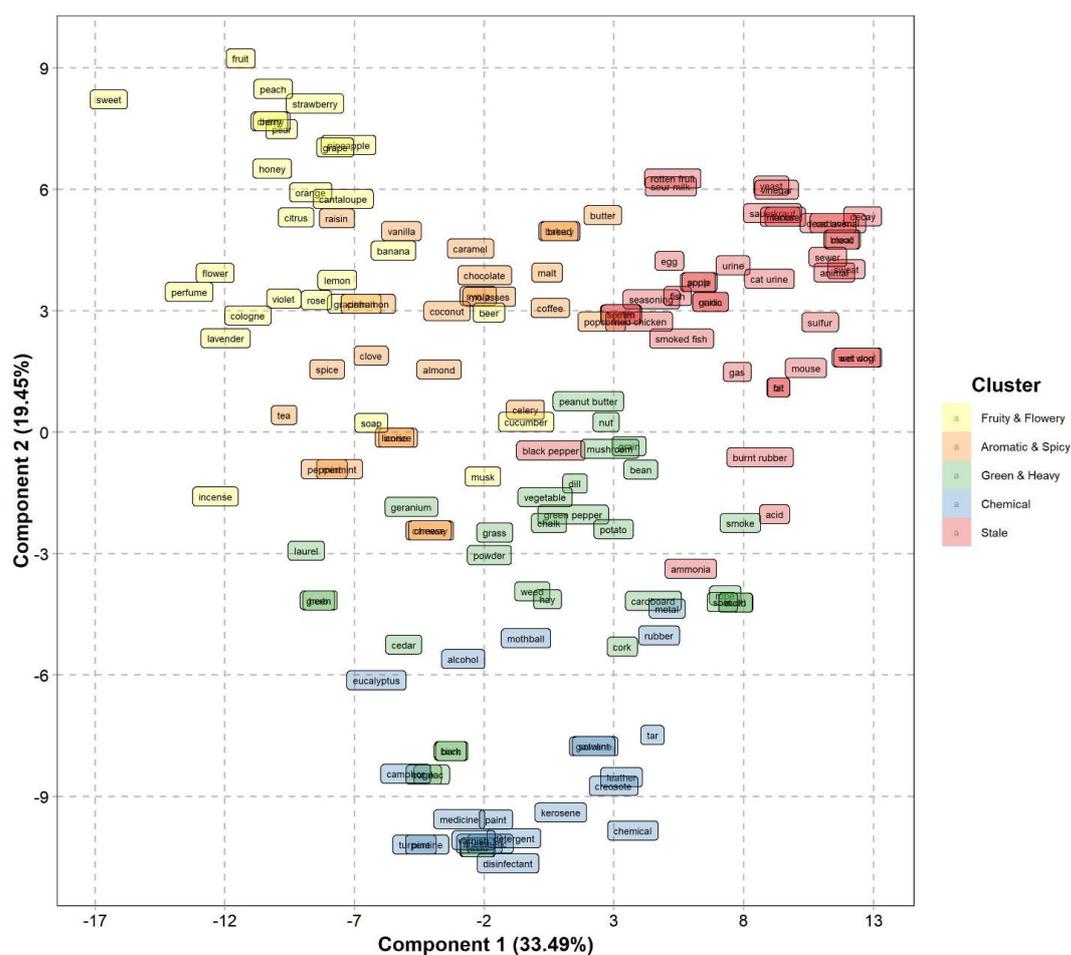
Figure 15. PCA score plot of PC1 and PC2 scores of the Dravnieks descriptors in the ratings-based PCA model, together with descriptor clusters identified in the hierarchical clustering analysis.

Table 6. Correlations between PC1 and PC2 of the word embedding PCA model, and PC1 and PC2 of the ratings-based PCA model of. P-values are corrected for multiple comparisons on the basis of the method proposed by Holm (1979).

| | | Word embedding model | | | |
| | | PC1 | | PC2 | |
		r	p	r	p
Ratings-	**PC1**	0.519	<.0001	0.381	<.0001
based model	**PC2**	-0.390	<.0001	0.329	.0005

is related to pleasantness, it is also related to edibility, thus primarily differentiating between pleasant/edible and pleasant/unedible. PC2, on the other hand, has a moderate positive correlation to both PC1 and PC2 of the ratings-based model. Figure 16 and 17 further shows that although descriptors referring to edibles score low on PC2, and descriptors referring to unedibles tend to score high, there is also a tendency for pleasant odor sources (e.g., *camphor, musk, cedar*) to score low, and unpleasant sources (e.g., *soup, fried chicken, mouse*) to score high. Thus, PC2 also seems to conflate edibility and pleasantness so that it primarily differentiates pleasant/unedible and unpleasant/edible.
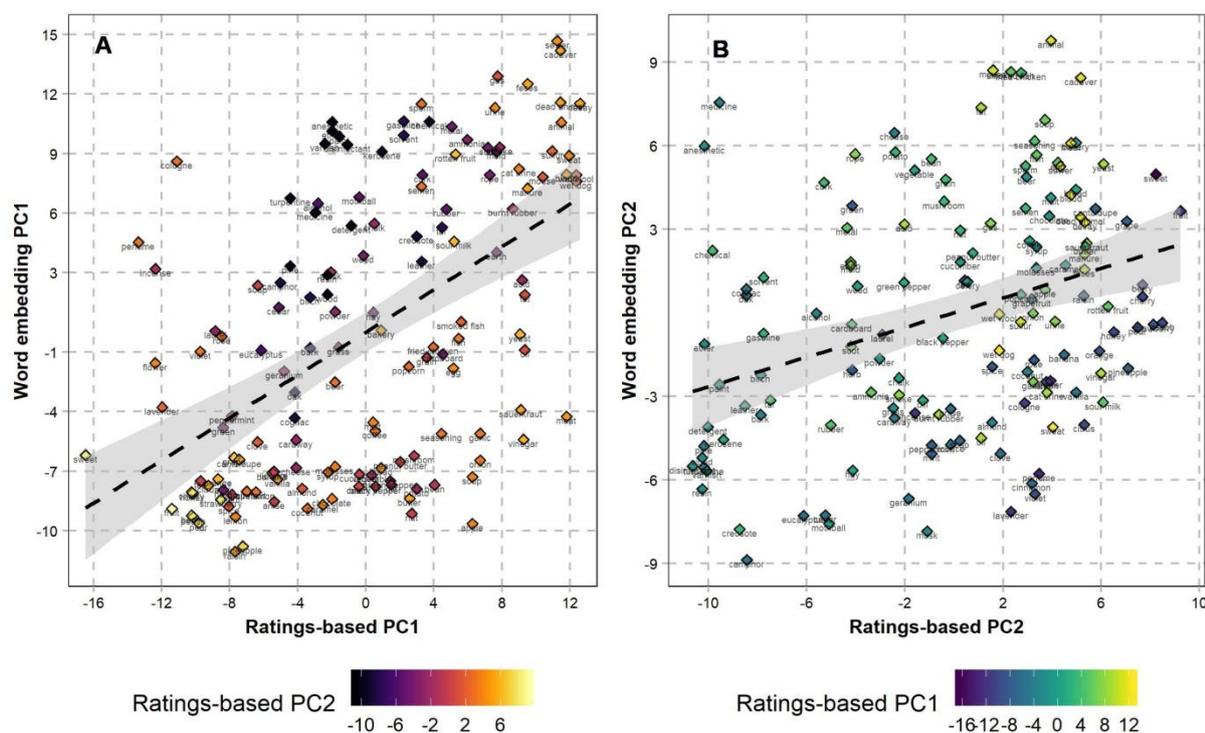

Figure 16. Correlations between the PC scores of the Dravnieks descriptor PCA models, comparing the PC scores of the word embedding model to the ratings-based model. A. The correlation between PC1 of the ratings-based model and PC1 of the word embedding model, also illustrating PC2 of the ratings-based model. B. The correlation between PC2 of the ratings-based model and PC2 of the word embedding model, also illustrating PC1 of the ratings-based model.
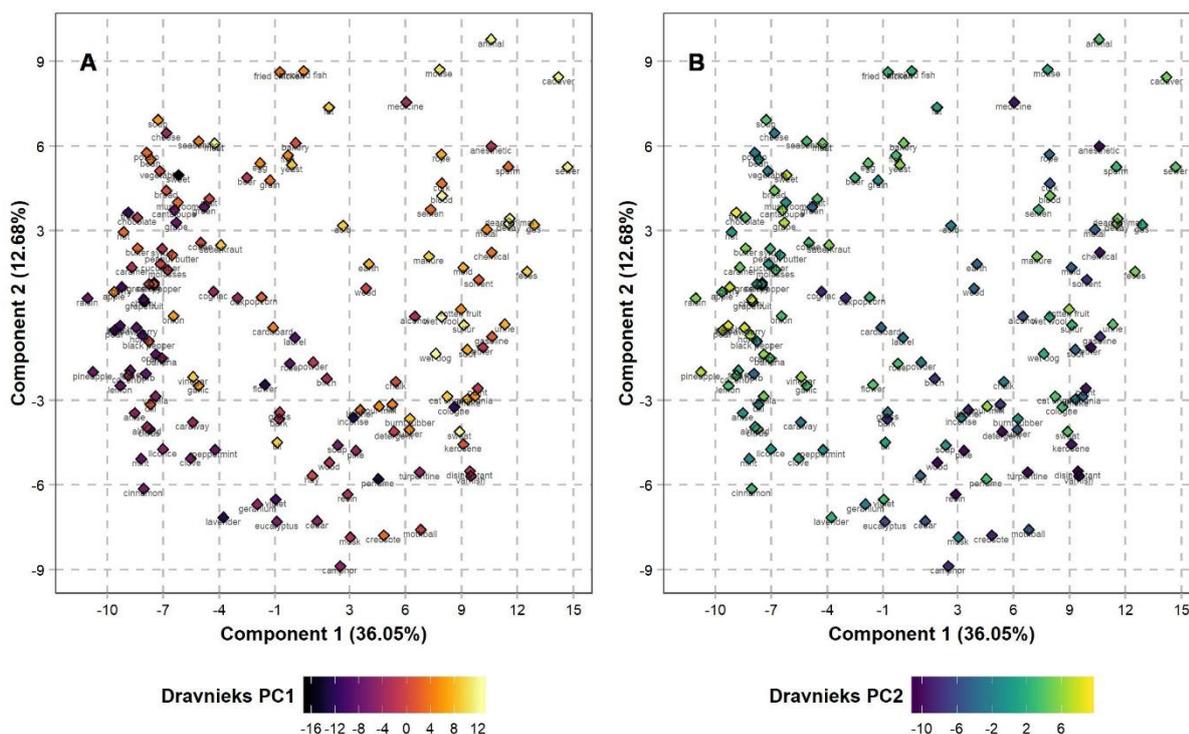
Figure 17. PCA score plots of PC1 and PC2 scores of the word embedding model of the Dravnieks descriptors, also showing PCA scores of the corresponding ratings-based model. A. PC1-PC2 score plot of the word embedding model, also showing PC1 scores of the ratings-based PCA model. B. PC1-PC2 score plot of the word embedding model, also showing PC1 scores of the ratings-based PCA model.

Discussion

In this paper, we develop and test a method that uses structured collections of natural texts and machine learning algorithms to automatically identify odor and flavor descriptors and derive the semantic organization of those descriptors. Our method automatically identifies descriptors in texts on the basis of their olfactory and gustatory association (i.e., OAI, see Iatropoulos et al., 2018), as based on their distributions in olfactory and gustatory contexts in natural text collections. The semantic space of these descriptors is then derived using a distributional-semantic word embedding model, which represents semantic distances between words as vector distances in a multi-dimensional space. Our method can be used for a potentially unlimited set of descriptors without requiring any perceptual data.

A few recent studies have instead investigated the semantic and perceptual space of odors using natural language data. Some studies have used language data to quantify descriptors with respect to how strongly they are associated with olfaction and gustation. Croijmans et al. (2019) used a text-based, computational approach on a collection of wine reviews in order to identify the vocabulary used to describe wine qualities. Iatropoulos et al. (2018) developed the method employed in the present study, the Olfactory Association Index (OAI), which quantifies words with respect to how strongly associated they are with olfactory and gustatory contexts. Other studies have used pre-trained distributional-semantic word embedding models to predict the applicability of odor descriptors to odor molecules (Gutiérrez et al., 2018; Nozaki & Nakamoto, 2018). Gutiérrez et al. (2018) found that semantic vector representations can be used to predict odor-descriptor ratings, thereby showing that there is a high correspondence between a perceptual odor space derived from odor-descriptor ratings, on the one hand, and a semantic space derived from a distributional-semantic word embedding model, on the other.

Taken together, these studies show that corpus-based, computational methods can be used to automatically identify olfactory and gustatory vocabulary, and, further, to derive a semantic space of those descriptors which is similar to a space that is derived from perceptual data.

However, previous studies have either been limited to a pre-determined set of odors and descriptors, and therefore not applied to the full range of descriptors used in everyday situations, or they have not been concerned with the semantic relationships between the descriptors, and therefore been uninformative regarding how those descriptors are organized semantically. For example, although Croijmans and colleagues were able to identify the set of terms that most frequently are used to describe wine qualities, their findings were uninformative regarding how those terms are semantically related. In their wine aroma wheel, the automatically identified terms were "manually classified" (Crojimans et al. 2019: 14). In this study, we overcome these shortcomings by first automatically identifying smell and taste descriptors in texts on the basis of their olfactory association, and then deriving the semantic space of those descriptors is using word embedding modeling specifically trained and olfactory and gustatory contexts.

We investigated the semantic organization of the descriptors with the highest olfactory association in terms of OAI, both by looking at the top 50 source-based, top 50 abstract and top 50 evaluative descriptors combined, and by investigating the top 150 source-based, the top 150 abstract, and the top 85 evaluative descriptors. All analyses involved hierarchical clustering and PCA analysis of a descriptor distance matrix that was derived from the semantic vectors of the descriptors at hand. Previous work that has investigated the semantic organization of odor descriptors on the basis of perceptual data (e.g., Dravnieks 1985) have found that odor descriptors primarily are differentiated in terms of pleasantness, but also with respect to edibility (e.g., Kahn et al. 2007; Zarzo 2008; Koulakov et al. 2011; Licon et al. 2018). Further, the distinction between source-based, abstract and evaluative descriptors is highly related to a difference in concreteness, the former descriptors being concrete and the latter abstract. We therefore further investigated whether and to what extent the principal dimensions identified in our analyses corresponded to pleasantness, edibility, and concreteness.

The analysis of the three classes of descriptors combined clearly indicates that our method is able to categorize and distinguish descriptors on the basis of qualitative differences of the perceptual experiences that they denote, rather than primarily differentiating between them based on functional or linguistic criteria. It could very well be that functionally distinct classes of descriptors would cluster together, such as differentiating between food products and food-related things, hygiene- and body-related concepts, substances and materials, flowers and plants, abstract properties, and evaluative concepts. However, although there is a tendency for descriptors from some of these classes to co-cluster, many of them are distributed across several of the identified clusters (see Figure 1). Food products and abstract properties are distributed across all clusters. Evaluative concepts are found in four clusters. Hygiene- and body related concepts are found in two clusters, and so is also substances and materials, as well as flowers and plants. Further, given that the distance between the word vectors of the word embedding model is based on the syntactic context that the descriptors occur in, which in turn is co-dependent on other lexical properties such as word morphology and part-of-speech, descriptors could also be expected to cluster on the basis of non-semantic lexical properties. For instance, the clustering could primarily differentiate between nouns and adjectives, or source-based descriptors from non-source based descriptors, which would further be divided into abstract and evaluative descriptors. But this is not the case either - nouns and adjectives are represented in all five clusters, and all three descriptor types are found in four out of five clusters. The clustering is instead primarily done on the basis of pleasantness, with descriptors referring to unpleasant experiences in the first cluster and descriptors denoting pleasant experiences in the last. Edibility also seems to play it's part. The two main partitionings, differentiating between pleasant and unpleasant experiences, are further divided into clusters which either mainly contain descriptors associated with non-edibles (i.e.., the clusters *Stale & Repulsive* and *Floral & Fragrant*) or with edibles (i.e., *Distasteful & Unpleasant*, *Spicy & Sweet* and *Distinct & Appetizing*), thereby also distinguishing descriptors in terms of edibility.

A similar conclusion can be drawn from the PCA analysis. Here, the primary dimension seems to differentiate between source-based and abstract descriptors referring to pleasant edibles or pleasant flavor and odor qualities, on the one hand, and abstract and evaluative descriptors referring to more general unpleasant odor qualities, on the other. Thus, the PCA analysis primarily differentiates descriptors on the basis of valence / pleasantness, but edibility is also of importance, and these dimensions seems to be confounded to some extent. The second dimension of the PCA analysis seems to be related to olfactory association, differentiating between descriptors that most frequently are used to express odor experiences, from those that more often also are used in other situations. The third dimension is primarily related to a concreteness cline more directly, differentiating between highly abstract and evaluative descriptors, on the one hand, and highly concrete, source based-descriptors, on the other, independent of pleasantness and edibility. Here, the importance of non-semantic lexical differences between descriptors is shown.

The analyses of descriptors within each descriptor category further indicate that or method primarily differentiate descriptors on the basis of pleasantness and edibility, although these dimensions seem to be weighted differentially between the three classes. For source-based descriptors, the hierarchical clustering primarily differentiates descriptors referring to edibles from descriptors denoting non-edibles, further dividing those on the basis of pleasantness. However, in the PCA analysis, descriptors are primarily differentiated in terms of pleasantness, and secondarily with respect to edibility. For abstract descriptors, the opposite pattern is found. The hierarchical clustering primarily differentiates abstract descriptors in terms of pleasantness, whereas in the PCA analysis, the primary dimension seems to be related to edibility, and the secondary dimension to pleasantness. Evaluative descriptors, finally, are almost exclusively differentiated on the basis of pleasantness both in the hierarchical clustering and the PCA analysis. Given the semantic nature of this class of descriptors, differentiating between objects, events and experiences in terms of affective quality, this is to be expected and provides further support for the feasibility of our method.

As a final step, we compared a word embedding-based descriptor space to a ratings-based descriptor space, derived from the descriptors ratings in the Dravnieks data set (Dravnieks, 1985). This was done by comparing the fit between the descriptor clustering and descriptor PCA space based on the semantic vectors, on the one hand, to the ratings-based descriptor clustering and PCA space. The comparison between the clustering solutions showed that whereas the ratings-based clustering primarily differentiate between descriptors with respect to pleasantness, the word embedding clustering primarily distinguish descriptors on the basis of edibility. However, qualitatively, there is a great deal of similarities between the five main clusters of each cluster solution, and our analyses of the cophenetic correlation (Sokal & Rohlf, 1962) between the clusterings show that the overlap is far greater than would be expected under the null hypothesis of no overlap. In our comparison of the PCA spaces, we found the primary dimension of the word embedding PCA space to primarily be related to the primary dimension of the ratings-based space, but also to the secondary dimension of the ratings-based space, thus differentiating between pleasant/edible and pleasant/unedible. The second dimension of the word embedding space was moderately related to both the primary and secondary dimension of the ratings-based space, distinguishing between pleasant/unedible and unpleasant/edible. Thus, whereas a descriptor space derived from odor-descriptor ratings keep the pleasantness and the edibility dimensions orthogonal, a descriptor space derived from word embeddings will to a greater extent confound these dimensions.

We have developed a method that automatically identifies odor and flavor descriptors in natural texts on the basis of how probable they are to occur in olfactory and gustatory contexts, as opposed to other contexts. Our method then derives the semantic organization of the identified descriptors on the basis of their distribution in those contexts. As such, an important part of the method involves identifying olfactory and gustatory contexts, which is done by identifying a set of olfactory- or gustatory-related key words such as 'odor' and 'taste'. It is an open question as to how robust our method is with respect to which key words are included. It could very well be that it is highly sensitive to the key words that

are being used, and that our results would have been quite different had we chosen a smaller set of key words, or limited ourselves to olfactory words. This could be further evaluated in future research.

Conclusions

In this paper, we have presented and evaluated a method that uses a fully data-driven and automatic approach to identify olfactory and gustatory descriptors and derive their semantic organization, using natural language data. In this method, descriptors are identified on the basis of their olfactory and gustatory association. Their semantic organization is then derived from a distributional-semantic word embedding model, that is specifically trained on olfactory and gustatory contexts. Our method is able to categorize and distinguish descriptors on the basis of qualitative differences of the perceptual experiences that they denote, in that it primarily differentiates between descriptors in terms of pleasantness / valence and edibility. The semantic space derived by our method is further similar to a space that is derived from perceptual data, such as odor-descriptor ratings. However, whereas the pleasantness and edibility dimensions are kept orthogonal in a semantic space derived from perceptual data, these two dimensions are to a greater extent confounded in semantic space derived with our method. Further work is needed to further assess the robustness of our method.

References

Bird, S., Klein, E., & Loper, E. (2009). *Natural language processing with Python* (1st ed). O'Reilly.

Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, *46*(3), 904–911. https://doi.org/10.3758/s13428-013-0403-5

Burenhult, N., & Majid, A. (2011). Olfaction in Aslian Ideology and Language. *The Senses and Society*, *6*(1), 19–29. https://doi.org/10.2752/174589311X12893982233597

Burlingame, G. A., Suffet, I. H., Khiari, D., & Bruchet, A. L. (2004). Development of an odor wheel classification scheme for wastewater. *Water Science and Technology*, *49*(9), 201–209. https://doi.org/10.2166/wst.2004.0571

Castro, J. B., Ramanathan, A., & Chennubhotla, C. S. (2013). Categorical Dimensions of Human Odor Descriptor Space Revealed by Non-Negative Matrix Factorization. *PLoS ONE*, *8*(9), e73289. https://doi.org/10.1371/journal.pone.0073289

Croijmans, I., Hendrickx, I., Lefever, E., Majid, A., & Van Den Bosch, A. (2019). Uncovering the language of wine experts. *Natural Language Engineering*, 1–20. https://doi.org/10.1017/S1351324919000500

Croijmans, I., & Majid, A. (2016). Not All Flavor Expertise Is Equal: The Language of Wine and Coffee Experts. *PLOS ONE*, *11*(6), e0155845. https://doi.org/10.1371/journal.pone.0155845

De Wijk, R. A., & Cain, W. S. (1994). Odor quality: Discrimination versus free and cued identification. *Perception & Psychophysics*, *56*(1), 12–18. https://doi.org/10.3758/BF03211686

Dravnieks, A. (1985). *Atlas of odor character profiles*. ASTM. http://www.astm.org/DIGITAL_LIBRARY/MNL/SOURCE_PAGES/DS61.htm

Dunn, J. C. (1974). Well-Separated Clusters and Optimal Fuzzy Partitions. *Journal of Cybernetics*, *4*(1), 95–104. https://doi.org/10.1080/01969727408546059

Engen, T. (1991). *Odor sensation and memory*. Praeger Publishers.

Fisher, R. M., Barczak, R. J., Suffet, I. H. "Mel," Hayes, J. E., & Stuetz, R. M. (2018). Framework for the use of odour wheels to manage odours throughout wastewater biosolids processing. *Science of The Total Environment*, *634*, 214–223. https://doi.org/10.1016/j.scitotenv.2018.03.352

Gutiérrez, E. D., Dhurandhar, A., Keller, A., Meyer, P., & Cecchi, G. A. (2018). Predicting natural language descriptions of mono-molecular odorants. *Nature Communications*, *9*(1), 4979. https://doi.org/10.1038/s41467-018-07439-9

Han, L., Kashyap, A. L., Finin, T., Mayfield, J., & Weese, J. (2013). *UMBC_EBIQUITY-CORE: Semantic Textual Similarity Systems*. 44–52.

Herz, R. S., & Engen, T. (1996). Odor memory: Review and analysis. *Psychonomic Bulletin & Review*, *3*(3), 300–313. https://doi.org/10.3758/BF03210754

Holm, S. (1979). A Simple Sequential Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, *6*, 65–70.

Iatropoulos, G., Herman, P., Lansner, A., Karlgren, J., Larsson, M., & Olofsson, J. K. (2018). The language of smell: Connecting linguistic and psychophysical properties of odor descriptors. *Cognition*, *178*, 37–49. https://doi.org/10.1016/j.cognition.2018.05.007

Jellinek, J. S. (1997). The psychological basis of perfumery: A re-evaluation. In J. S. Jellinek (Ed.), *The Psychological Basis of Perfumery* (pp. 240–257). Springer Netherlands. https://doi.org/10.1007/978-94-009-1567-1_26

Jönsson, F. U., & Stevenson, R. J. (2014). Odor Knowledge, Odor Naming, and the "Tip-of-the-Nose" Experience. In B. L. Schwartz & A. S. Brown (Eds.), *Tip-of-the-Tongue States and Related Phenomena* (pp. 305–326). Cambridge University Press. https://doi.org/10.1017/CBO9781139547383.014

Kaeppler, K., & Mueller, F. (2013). Odor Classification: A Review of Factors Influencing Perception-Based Odor Arrangements. *Chemical Senses*, *38*(3), 189–209. https://doi.org/10.1093/chemse/bjs141

Kaufman, L., & Rousseeuw, P. J. (Eds.). (1990). *Finding Groups in Data*. John Wiley & Sons, Inc. https://doi.org/10.1002/9780470316801

Keller, A., Gerkin, R. C., Guan, Y., Dhurandhar, A., Turu, G., Szalai, B., Mainland, J. D., Ihara, Y., Yu, C. W., Wolfinger, R., Vens, C., Schietgat, L., De Grave, K., Norel, R., DREAM Olfaction Prediction Consortium, Stolovitzky, G., Cecchi, G. A., Vosshall, L. B., & Meyer, P. (2017). Predicting human olfactory perception from chemical features of odor molecules. *Science*, *355*(6327), 820–826. https://doi.org/10.1126/science.aal2014

Khan, R. M., Luk, C.-H., Flinker, A., Aggarwal, A., Lapid, H., Haddad, R., & Sobel, N. (2007). Predicting Odor Pleasantness from Odorant Structure: Pleasantness as a Reflection of the Physical World. *Journal of Neuroscience*, *27*(37), 10015–10023. https://doi.org/10.1523/JNEUROSCI.1158-07.2007

Koulakov, A. A. (2011). In search of the structure of human olfactory space. *Frontiers in Systems Neuroscience*, *5*. https://doi.org/10.3389/fnsys.2011.00065

Licon, C. C., Manesse, C., Dantec, M., Fournel, A., & Bensafi, M. (2018). Pleasantness and trigeminal sensations as salient dimensions in organizing the semantic and physiological spaces of odors. *Scientific Reports*, *8*(1), 8444. https://doi.org/10.1038/s41598-018-26510-5

Lynott, D., & Connell, L. (2009). Modality exclusivity norms for 423 object properties. *Behavior Research Methods*, *41*(2), 558–564. https://doi.org/10.3758/BRM.41.2.558

Lynott, D., & Connell, L. (2013). Modality exclusivity norms for 400 nouns: The relationship between perceptual experience and surface word form. *Behavior Research Methods*, *45*(2), 516–526. https://doi.org/10.3758/s13428-012-0267-0

Maechler, M., Rousseeuw, P., Struyf, A., Hubert, M., & Hornik, K. (2019). *cluster: Cluster Analysis Basics and Extensions* (Version R package version 2.1.0) [R].

Majid, A., & Burenhult, N. (2014). Odors are expressible in language, as long as you speak the right language. *Cognition*, *130*(2), 266–270. https://doi.org/10.1016/j.cognition.2013.11.004

Majid, A., Roberts, S. G., Cilissen, L., Emmorey, K., Nicodemus, B., O'Grady, L., Woll, B., LeLan, B., de Sousa, H., Cansler, B. L., Shayan, S., de Vos, C., Senft, G., Enfield, N. J., Razak, R. A., Fedden, S., Tufvesson, S., Dingemanse, M., Ozturk, O., … Levinson, S. C. (2018). Differential coding of perception in the world's languages. *Proceedings of the National Academy of Sciences*, *115*(45), 11369–11376. https://doi.org/10.1073/pnas.1720419115

Marcus, M., Kim, G., Marcinkiewicz, M. A., MacIntyre, R., Bies, A., Ferguson, M., Katz, K., & Schasberger, B. (1994). *The Penn Treebank: Annotating Predicate-Argument Structure* [Unpublished Work].

Noble, A. C., Arnold, R. A., Buechsenstein, J., Leach, E. J., Schmidt, J. O., & Stern, P. M. (1987). Modification of a Standardized System of Wine Aroma Terminology. *American Journal of Enology and Viticulture*, *38*(2), 143–146.

Nozaki, Y., & Nakamoto, T. (2018). Predictive modeling for odor character of a chemical using machine learning combined with natural language processing. *PLOS ONE*, *13*(12), e0208962. https://doi.org/10.1371/journal.pone.0208962

Olofsson, J. K., & Gottfried, J. A. (2015). The muted sense: Neurocognitive limitations of olfactory language. *Trends in Cognitive Sciences*, *19*(6), 314–321. https://doi.org/10.1016/j.tics.2015.04.007

{R Core Development Team}. (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. http://www.R-project.org/

Řehůřek, R., & Sojka, P. (2010). Software Framework for Topic Modelling with Large Corpora. *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, 45–50.

Sokal, R. R., & Rohlf, F. J. (1962). THE COMPARISON OF DENDROGRAMS BY OBJECTIVE METHODS. *TAXON*, *11*(2), 33–40. https://doi.org/10.2307/1217208

Suffet, I. H., & Rosenfeld, P. (2007). The anatomy of odour wheels for odours of drinking water, wastewater, compost and the urban environment. *Water Science and Technology*, *55*(5), 335–344. https://doi.org/10.2166/wst.2007.196

Taylor, A., Marcus, M., & Santorini, B. (2003). The Penn Treebank: An overview. In A. Abillé (Ed.), *Treebanks—Building and Using Parsed Corpora* (pp. 5–22). Kluwer Academic Publishers.

Ward, J. H. (1963). Hierarchical Grouping to Optimize an Objective Function. *Journal of the American Statistical Association*, *58*(301), 236–244. https://doi.org/10.1080/01621459.1963.10500845

Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods*, *45*(4), 1191–1207. https://doi.org/10.3758/s13428-012-0314-x

Wnuk, E., & Majid, A. (2014). Revisiting the limits of language: The odor lexicon of Maniq. *Cognition*, *131*(1), 125–138. https://doi.org/10.1016/j.cognition.2013.12.008

Zarzo, M. (2008). Psychologic Dimensions in the Perception of Everyday Odors: Pleasantness and Edibility. *JOURNAL OF SENSORY STUDIES*, *3*, 354.

Zarzo, M., & Stanton, D. T. (2009). Understanding the underlying dimensions in perfumers' odor perception space as a basis for developing meaningful odor maps. *Attention, Perception & Psychophysics*, *71*(2), 225–247. https://doi.org/10.3758/APP.71.2.225