# Integrating Heterogeneous Odor Response Data into a Common Response Model: A DoOR to the Complete Olfactome

## C. Giovanni Galizia[1], Daniel Münch[1], Martin Strauch[1], Anja Nissler[2] and Shouwen Ma[1]

[1]Department of Neurobiology, University of Konstanz, 78457 Konstanz, Germany and
[2]Department of Biology, Humboldt University of Berlin, Berlin 10099, Germany

Correspondence to be sent to: C. Giovanni Galizia, Department of Neurobiology, University of Konstanz, 78457 Konstanz, Germany. e-mail: giovanni.galizia@uni-konstanz.de

## Abstract

We have developed a new computational framework for merging odor response data sets from heterogeneous studies, creating a consensus metadatabase, the database of odor responses (DoOR). As a result, we obtained a functional atlas of all available odor responses in Drosophila melanogaster. Both the program and the data set are freely accessible and downloadable on the Internet (http://neuro.uni-konstanz.de/DoOR). The procedure can be adapted to other species, thus creating a family of "olfactomes" in the near future. Drosophila melanogaster was chosen because of all species this one is closest to having the complete olfactome characterized, with the highest number of deorphanized receptors available. The database guarantees long-term stability (by offering time-stamped, downloadable versions), up-to-date accuracy (by including new data sets as soon as they are published), and portability (for other species). We hope that this comprehensive repository of odor response profiles will be useful to the olfactory community and to computational neuroscientists alike.

Key words: computational model, Drosophila melanogaster, metadatabase, odor responses, olfactory receptor, response profiles

## Introduction

The aim of neuroscience is to understand the brain based on empirically measured physiological data. The community, therefore, relies on access to good experimental data, and considerable effort is being made to create databases that offer large, annotated data sets from physiological experiments made across the world in many laboratories (Herz et al. 2008). However, a major difficulty lies in the comparability of data that come from different places and times. Small changes in experimental parameters can influence the outcome of a physiological experiment, and even under similar conditions, different groups might use other readout parameters for physiological activity. For example, stimulus response intensity might be reported in spike counts, spike rates, or calcium concentration changes.

Odors consist of volatile airborne molecules that can be perceived by an organism. In the olfactory system, odors are recognized by a large family of odor receptors (ORs). In most animals, including humans, mice, and the fruit fly Drosophila melanogaster, each receptor cell expresses one or a few receptor proteins, which give that cell a specific odor response profile. This profile can be represented by a function: to any given

chemical representing an odor stimulus we can map a response intensity. Because most chemicals will elicit responses in more than 1 receptor cell type, each odor elicits a combinatorial activity pattern across these channels. It is this combinatorial nature of olfaction that allows the brain to recognize and remember thousands or maybe millions of different odors with a limited number of receptor types: approximately 350 in humans (Glusman et al. 2001), 1000 in mice (Buck and Axel 1991), and 60 in D. melanogaster (Vosshall et al. 1999). In order to understand how the brain perceives an odor, the ideal situation would be to know all response profiles of all receptors for a given species. Because of technical difficulties, most receptor types are still orphans, that is, their ligands are unknown. The most prominent exception to this is the fruit fly D. melanogaster, where many studies have measured odor response patterns in individual cells and in small groups of cells, either in vivo or in vitro. These odor response profiles in D. melanogaster come from different research groups, which have used different techniques (e.g., heterologous expression, Smart et al. 2008; in situ recordings in wild-type sensilla, de Bruyne et al. 1999; in situ recordings in the "empty

neuron," Hallem et al. 2004; calcium imaging of cellular responses, Pelz et al. 2006). Furthermore, the set of tested odors differed across studies. As a consequence, it is difficult to compare different studies numerically. Yet, no study has covered all receptor cells so far, and given the resources needed for such an enterprise, it would appear as a waste to do so now, where many receptors have already been deorphanized in great detail.

Exploiting this wealth of data available from the fruit fly, we have therefore developed a new approach that allows us to compare and combine odor response profiles from many studies even when their physiological responses are heterogeneous due to different techniques used and when the odors tested are only partially overlapping. As a result, we obtain consensus profiles that are based on many studies and thus are statistically more reliable than any single study. We have developed a software platform that allows to extract odor response profiles across chemicals for individual receptors or to extract the entire combinatorial response pattern elicited by a given chemical. The software is open source and can be modified by the user. Although we will update the database on a regular basis, the database includes a feature that allows for retrieving the state of the database at any given time in the past. This is important to allow for comparative computational studies on reference data sets.

The database is suitable for further studies into the combinatorial nature of olfactory coding, into the logic of ligand receptor interaction in olfactory receptors, and for other applications. Furthermore, the software can be used to create similar databases for other species, including mice and humans, as soon as enough data will be available. Thus, it joins related efforts for databases of olfactory receptor sequences and their ligands (Crasto et al. 2002), as well as other data repositories, for example, http://senselab.med.yale.edu/senselab/ordb or http://gara.bio.uci.edu/. The database of odor responses (DoOR) package is available from http://neuro.uni-konstanz.de/DoOR.

## Materials and methods

### Nomenclature

Receptors (e.g., *dOr22a* and *Ir76b*), receptor cells (e.g., ab3A and ac3B), and corresponding glomeruli (e.g., DM2 and VC3l) were labeled following the standards in *D. melanogaster* literature (see Laissue et al. 1999 for glomerulus nomenclature). ORs in *D. melanogaster* belong to 3 major families: ORs, gustatory receptors, and ionotropic receptors (Larsson et al. 2004; Kwon et al. 2007; Benton et al. 2009). Each odor is given by its chemical name (e.g., 2-heptanone) and the unique Chemical Abstracts Service number (http://www.cas.org).

### Sources for published odor response profiles

Odor responses were taken from studies with at least 5 odors tested for a given receptor. Each study enters the database with its own name based on the author, the publication year, and a short data descriptor. For example, the data from Hallem (Hallem et al. 2004) enter the database as 2 data sets called Hallem.2004.EN and Hallem.2004.WT. Here, EN stands for an empty neuron recording, where receptor proteins are ectopically expressed in an empty olfactory neuron, whereas WT signifies a wild-type recording, that is, a recording from an olfactory neuron that naturally expresses its receptor protein. A list of all studies with nomenclature and details on the respective experiments is provided (Supplementary Table S2). As most studies reported only one odorant concentration level, no information about response properties across concentration ranges is included in the present version of the database.

### Sources for unpublished odor response profiles

We recorded odor response profiles for *dOr13a*, *dOr67b*, and *dOr92a*. We used OrXX:GAL4 and UAS:G-CaMP flies and recorded calcium responses using a CCD (charge-coupled device) camera and a ×50 air objective through the intact antenna cuticle as described in detail elsewhere (Pelz et al. 2006). Odors were diluted in mineral oil in decadic steps ($10^{-2}$, $10^{-3}$, ...), with 1:100 ($10^{-2}$) as the highest concentration, to measure complete odor response curves. Five milliliters of diluted odor was kept in sealed 20-ml vials filled with nitrogen, and 2-ml headspace was used for each stimulation. Odor delivery was automated using a headspace multisampler adapted from gas chromatography (CombiPAL, CTC analytics). For each odor stimulus, a train of 80 fluorescent frames was recorded, with a sampling rate of 4 frames per second. Odor stimuli were applied as 2 pulses, each 1-s long, at time points 6 and 9 s in each measurement. Bleach-corrected odor responses were converted into relative fluorescence changes as $\Delta F/F$, with $F$ being the background fluorescence before odor stimulation. For each measurement, odor response magnitude was quantified as the average calcium increase in $\Delta F/F$ during 4 s after first stimulus onset. Maximum response magnitude varies across animals, mostly due to difference in G-CaMP expression levels and cuticle pigmentation darkness. Before averaging across animals, responses were therefore normalized within each animal by setting the response to a reference stimulus to 1 and scaling all other responses accordingly. The reference odor was 3-octanol (589-98-0) for *dOr13a*, 1-hexanol (111-27-3) for *dOr67b*, and 2,3-butanedione (431-03-8) for *dOr92a*.

### Preprocessing of odor response profiles

We transformed all data sets where values decrease for better ligands (i.e., data reported as 50% effective concentration (EC50) values of odor dilution) by inverting their values in the database (e.g., in Pelz.2006.AntEC50 an EC50 value of –4.13 is coded as +4.13 in the database) in order to comply with our assumption that $R_1(a) < R_1(b) \Rightarrow R_2(a) < R_2(b)$ for all odors a,b (see Results). Before fitting an odor response vector,

its values were all scaled to the range [0, 1] in order to avoid unequal weighting of the 2 vectors in the fitting procedure.

**Finding the best-fitting function**

Take a data set of odor response profiles covering $o_A$ odors in $r_A$ receptors. We write this data set as a matrix (see Supplementary Figure S7). We have several such data sets from different studies, and each study may cover a different (but overlapping) set of odors and a different (but overlapping) set of receptors. Let there be $s$ such studies, and let us denote them $A^1, \ldots A^s$. Thus, the response to odor $i$ in receptor $j$ for study $k$ is $A_{ij}^k$. For better readability, where useful, we denote columns by the corresponding receptor names and omit subscripts where the entire range is intended. Thus, $A_{[Or22a]}^k$ contains the column of odor responses for receptor 22a in the $k$th study. We will follow the *Or22a* example throughout this section. The goal of the algorithm is to merge all available $A^k$ in order to obtain a single consensus matrix $M \in \mathbb{R}^{r \times o}$, where $r$ is the number of all receptors and $o$ is the number of all odors. Merging is done sequentially for each receptor, and within each receptor, merging is done iteratively (Supplementary Figure S7). First, 2 data sets are merged and then the resulting consensus data set is merged to the next original data set. For small $s$ ($s$ may differ for different receptors), all possible merging sequences can be calculated. For large $s$, this exhaustive approach is not possible due to computing time constraints, and we follow a heuristic instead (see below).

For each merging step, we first fit 5 different monotonic functions to the pairs of data sets. The functions used are linear, exponential, sigmoid, asymptotic, and asymptotic with an offset (see Supplementary Figure S1 and user manual on the DoOR homepage). Fitting is done using the R routine nls(). This routine minimizes the square distance of the dependent variable $f(x)$ against the independent variable $x$. Graphically this corresponds to the vertical distances from each point onto that function. However, this is not the optimal solution because there is no "dependent" and "independent" data set. The best solution would be to minimize not the vertical distances but the perpendicular projections onto the fitted function. However, there is no efficient algorithm yet to do this calculation. Until such an algorithm will be implemented, we have taken an alternative approach: all 5 functions are also fitted flipping the 2 data sets, effectively optimizing not the vertical projections on the fit but the horizontal projections. In our algorithm, these are the "inverse" functions, so that effectively a total of 10 fitting functions were tested.

For each of these 10 fits, we calculate the average orthogonal distance (unlike the fitting of best parameters, for a set of given parameters this statistic is easily computed). We select the fitting function $f_{best}(x)$ with the smallest average orthogonal distance (mean distance [MD]). This function is only well defined within the data range of the 2 odor response

vectors that have been fitted, and an extrapolation beyond that range would create unwarranted results. Therefore, for values outside this range, we expand the function with a linear function, $f(x) = x + \text{intercept}$, where intercept is chosen to create a continuous function. Thus, the complete $f_{best}(x)$ consists of a linear function to the left, a fitted function in the center, and a linear function to the right.

**Merging 2 data sets**

For all odors present in both studies to be merged (or the study to be merged into the consensus set), the location of that odor on the trajectory of $f_{best}(x)$ is calculated by orthogonal projection. All odors that are present in only one of the 2 studies are also projected onto the function.

The odor response values of the newly merged set are calculated by measuring the distances along $f_{best}(x)$. Specifically, given a data point $p_1 = (x_1, y_1)$, we compute the distance from $p_{min} = (x_{min}, y_{min})$ to $p_1$ as follows:

$$d(p_{min}, p_1) = \int_{x_{min}}^{x_1} \sqrt{1 + \left(f'_{best}(x)^2\right)}\, dx.$$

This step is followed by scaling the whole range to [0, 1]. Now the complete data set, for this receptor, has 1 study less, and the procedure is iterated (Supplementary Figure S7).

**Data set merging order and data set exclusion**

When the number of data sets to be merged is large, not all merging orders can be tested. In this case, we first calculate merging quality (in terms of mean orthogonal distance) for all possible pairs and merge the 2 data sets that yield the best merging quality. This procedure is iterated until all data sets have been matched.

There are cases where no match is possible, and these data sets are excluded. First, the minimum overlap requested (in terms of common odors of both studies) is 4. Fewer overlapping odors do not give sufficient degrees of freedom to fit the monotonic functions. Second, only pairs that result in a mean orthogonal distance below 0.1415 (which corresponds to 10% of the maximum possible distance) are merged.

**Global scaling**

For comparison of responses across receptors (see Figure 3b), we developed a global scaling introducing a weighting factor $w_j$ for each receptor, making use of the information in studies that contain more than 1 receptor. Because studies that include many odors and receptors contain more across-receptor information, they are weighted more. Thus, for a study $k$, let $n.rec_k$ be the number of receptors covered and $n.odo_k$ the number of odors recorded. For each receptor $j$, in that study, we calculate $R_j^k$ as the maximum odor response within that receptor, and for that study, $S^k$ is the

maximum odor response across all receptors (in the units of that study, e.g., spikes per second). We then calculate:

$$w_j = \frac{\sum\limits_{k=1}^{s} n.rec_k \cdot \frac{R_j^k}{S^k} + \sum\limits_{k=1}^{s} n.odo_k \cdot \frac{R_j^k}{S^k}}{\sum\limits_{k=1}^{s} n.rec_k + \sum\limits_{k=1}^{s} n.odo_k}.$$

### Implementation and availability

All methods used in this work are implemented in the open source statistical environment R (R Development Core Team 2009). Apart from the source codes, the DoOR packages for R comprise the original data sets and a precomputed model response matrix. With a few R commands, the user can add data, compute his or her own model response matrix, and reproduce the plots from this paper. R can be obtained from www.r-project.org. The DoOR package is available from http://neuro.uni-konstanz.de/DoOR. A help file with detailed instructions can also be downloaded from that site.

For users who just wish to query the database without using the R package, we provide a web interface for the latest version of the database including 2D and 3D visualizations of the response patterns at http://neuro.uni-konstanz.de/DoOR.

## Results

### Fitting 2 data sets onto each other

Different odor response profile data sets can have very different qualities and data ranges. For example, studies reporting spike counts may have discrete values, for example, ranging from 0 to 500 spikes per second. Data based on calcium imaging may have percentage of fluorescence change values ranging from negative values (for inhibitory responses) to positive values (e.g., –5 to +18 $\Delta F/F$). Measurements that report receptor sensitivities calculated from entire dose response curves report data as the effective odor concentration that elicit half-maximal responses (EC50), with values ranging from, say, –6.0 to –2.0 (corresponding to log-based odor dilutions). Unlike the first 2 cases, better ligands have a lower value when expressed as EC50. With this heterogeneity in the qualitative nature of different data types, how could we combine them? Which is the property of odor response profiles that is, in theory, consistent across all data sets? We start with the observation that all odor response profiles of a particular receptor must be based on the same monotonic relationship. Given 2 odors a and b, we denote their responses with method 1 as $R_1(a)$ and $R_1(b)$ and with method 2 as $R_2(a)$ and $R_2(b)$. Our postulate states that $R_1(a) < R_1(b) \Rightarrow R_2(a) < R_2(b)$ for all a,b of a given odor response profile. Because all measurements have noise, this postulate will not be true in all real data sets, but the basic principle is that a better ligand in 1 data set should also be a better ligand in another data set.

We mapped data sets onto each other as pairs. In order to avoid too many free parameters, we selected 5 possible fitting models and their inverse (see Materials and methods): a linear model, an exponential, a sigmoid model, and 2 types of asymptotic nonlinear functions, 1 with an offset and 1 without (see Supplementary Figure S1). We show the merging of 2 data sets for dOr22a in Figure 1. This receptor has a broad response pattern, that is, many chemicals elicit responses (Figure 1a). Responses are plotted against each other for all odors that were measured in both sets (Figure 1b); note that values in Pelz.2006.AntEC50 range from 2 to 7 (negative logarithm of odor dilution necessary to elicit the half-maximal response), whereas responses in Hallem.2006.EN range from 0 to 250 (these are response frequencies in spikes per second, compare with Figure 1d). Different dimensionalities along the axes influence the fitting procedure (e.g., deviation along the spike axis would weigh more because the value ranges are larger). Therefore, each data set was linearly scaled to a common range [0, 1] before mapping (compare the axes in Figure 1b and c). A clear monotonic relationship (plus noise) is apparent between the 2 data sets.

Next, we mapped each point onto the regression function (Figure 1c). Because in these regressions both data sets are equal (i.e., there is no dependent variable), mapping is done by perpendicular projection, that is, we projected each data point onto the closest point on the regression function. Some odors were measured only in 1 of the 2 data sets. These odors were also projected onto the regression line. We did not extrapolate the fitting function beyond the data range covered by the 2 data sets. Rather, we projected values outside this range onto a unitary line (45° slope), thus leaving that range of the data set unaltered. Finally, we gave each point on the regression a value by calculating its position on the curve, scaled to the range [0, 1]. The resulting odor response profile was not the average of the 2 data sets but a fitted consensus set (Figure 1d). A comparison of the consensus set with the 2 original sets showed a good correspondence but also showed that for some odors the information in 1 set differed from the information in the other set. In no case, we attempted to weigh data sets based on our judgment of their quality: the more data sets are integrated the more individual outliers should become irrelevant.

Note that scaling to the [0, 1] interval might cause problems, for example, in case of a data set consisting only of weak ligands when compared with a data set with mostly strong ligands or when several receptors are compared. The first problem is addressed by not extrapolating the fitting function but using a unitary line beyond the range of each study. For the second case, we employed a global scaling to enable across-receptor comparisons (see Materials and methods and below).

**Figure 1**    Merging 2 response data sets for 1 receptor. **(a)** Tuning breadth of odor response profiles for *dOr22a* taken from 2 published sets: Pelz.2006.AntEC50 (Pelz et al. 2006) (top, ordinate units are percentage of calcium responses) and Hallem.2006.EN (Hallem and Carlson 2006) (bottom, ordinate values are spikes per second). Responses are arranged with strongest odor at the center in order to show the broad odor response profile confirmed in both studies, irrespective of the recording technique. Pelz reported EC50 values based on calcium responses from dose response profiles. Hallem reported action potential frequencies in the empty neuron preparation. **(b)** Plotting odor responses to substances that were measured in both data sets against each other shows a strong correspondence. Note that the values differ: spikes range (abscissa) from approximately 0 to 250 (spontaneous rate was not subtracted) and EC50 (ordinate) ranges from 1 to 7 (negative logarithm of odor dilution). **(c)** Generation of a consensus data set. Vertical projections from the circles as in (b) to a fitted regression function yield the consensus odor response. Odors that were measured only in 1 study are projected from the respective axis onto the regression curve (blue lines for Pelz and yellow lines for Hallem). Consensus responses are calculated from the position along the regression curve. **(d)** Comparison of odor response profiles of the overlapping odor set for the model responses and the 2 original data sets (EC50 and spikes per second, respectively). The model responses were arranged in decreasing order, whereas the other 2 data sets were ordered by matching the odors to the model response plot. The model response covers the normalized range [0, 1].

## Merging multiple data sets

Ideally, each receptor has been recorded in several studies giving rise to several data sets, with many overlapping odor responses. Merging data sets was done by iteration. To this end, we performed pairwise data set mapping with each of the fitting functions, and the function with the fit performance (lowest "MD") was noted. This results in a fit-quality matrix of all data sets, from which a cluster dendrogram can be derived for visualization when fit quality is interpreted as

similarity (Figure 2a). Note that this data set is also influenced by how many odors overlap between 2 data sets. In the extreme case, 2 sets with an overlap of just 2 odors would have a perfect fit even though they would not share any information about the odor response profile. Therefore, to create the dendrogram, we did only use those pairs that had at least 4 common values.

Next, the pair with the best-fit performance was merged. In Figure 2a, this corresponds to joining the 2 data sets with the highest node. As a result, the complete data contained 1 set less altogether. In the next step, the created merged set was taken as reference, and its fit performance with all other data sets was measured (Figure 2b). The data set with the lowest MD was merged into the reference, and this procedure was iterated until either all sets were merged into the consensus set or the breakout criterion was reached (see Materials and methods). With increasing number of studies, the reference set contains an increasing number of odor responses. Figure 2c shows the whole procedure for *dOr22a*, which is the receptor for which most studies were available. Because the sequence of merging studies slightly influences the outcome of the consensus data set, in cases where computationally feasible, we merged the data calculating all possible merging sequences and selected the best sequence on the basis of the mean deviation of the merged sequence to each original data set.

## Validation and rescaling

As a result, we obtained a consensus odor response profile as shown for a subset of odors with *dOr22a* in Figure 2d. How reliable are the individual values? We ran the merging process as many times as there were data sets, with each time 1 data set being dropped from the list. Therefore, for each odor, we obtained several data points, that is, as many as the number of studies that covered that odor and obtained error bars as shown in Figure 2d. These error bars confirmed that our approach yields reliable values.

Although remapping of odor responses to [0, 1] is useful for theoretical analysis of olfactory coding, in an experimental setting, odor responses are more useful if they are given in the same unit as the experiments themselves. Therefore, the package can be used to back project the merged data set onto the experimental data sets. Most importantly, the back-projected data set contained data points that were not measured in the original study but that can be directly compared with their numerical value (see Supplementary Figure S2).

SFR denotes "spontaneous firing rate," which is not an odor response but background activity in the absence of a stimulus. If upon stimulation with an odor firing rate drops below SFR this indicates an inhibitory response. Not all studies reported the SFR value, and some techniques have no access to this value. For example, calcium-imaging studies cannot measure uniform spontaneous activity (bursty spontaneous activity can be measured, Galan et al. 2006). In calcium-imaging studies, however, inhibitory responses

**Figure 2** Mapping many response sets for 1 receptor. **(a)** Hierarchical cluster dendrogram based on best-fit values of 10 data sets from 8 studies (de Bruyne et al. 2001; Dobritsa et al. 2003; Stensmyr et al. 2003; Hallem et al. 2004; Pelz 2005; Hallem and Carlson 2006; Pelz et al. 2006; Schmuker et al. 2007) with odor responses for *dOr22a*. The 2 sets with the best pairwise fit are Dobritsa.2003.EN and Schmuker.2007.TR. These 2 sets are

are visible as calcium concentration decreases, as opposed to the responses to control, air or mineral oil, which generally give no responses. In our procedure, as explained so far, the merged data were scaled to the range [0, 1]. SFR, air, and solvent were always treated as if they were stimuli, and thus, inhibitory responses could be recognized as values smaller than the SFR value. However, this is not always satisfactory, in particular when comparing different receptors that might have different levels for SFR. Therefore, data can be linearly rescaled to have the range SFR to maximum map into the range [0,1], and negative values as large as dictated by the linear fitting.

### Comparisons across receptors

Up to this point, all procedures were applied to each receptor per se without any comparison to responses in other receptors. Tuning breadth displays for 6 different receptors are shown in Figure 3a: for example, *Or67a* had a broad response profile, whereas *Or59b* had a sharp response profile. Note also that for some receptors, only few odor responses were known (e.g., *Or59c*). For each receptor, the maximum response was set to 1 and SFR was set to 0, making negative responses immediately visible.

However, the very nature of olfactory coding is combinatorial, and for the olfactory system as a whole, no response in a single receptor neuron type contains information without a comparison to other receptors (with the possible exception of very few labeled line systems). Assume, for example, that a receptor, *dOrX*, has so far only been measured with very weak ligands (i.e., no better ligand is as yet known). In this case, the procedure above would still give the best odor in the test set a value of 1, which when compared across receptors would be misleading. In order to compare receptors, it was therefore necessary to rescale them (see Materials and methods).

For the 6 receptors shown in Figure 3a, the rescaled results are shown in Figure 3b (see also Supplementary Figure S6).

then merged and create the first model response. **(b)** Best fit of the remaining 8 data sets with this modeled response (merged_data) shows that Bruyne.2001.WT is the next best match (smallest MD). This set is now merged with merged_data. This procedure is iterated for all sets that match merging criteria (see text). **(c)** Iterative sequence for *dOr22a* showing how for each step a different mapping function might be best. Here, Dobritsa.2003.EN is first merged to Schmuker.2007.TR (see a) using inv.sigmoid as function, yielding *merged_data[1]*. Each of the next frame gives the fitting function used, the number of odors common to both sets (*n*), and indicates new odors added into *merged_data[i+1]* by yellow vertical lines and odors present in *merged_data[i]* but not in the data set by blue horizontal lines. **(d)** Responses to 19 selected odors in *dOr22a*, as calculated from all available data sets. Ethyl hexanoate and methyl hexanoate are the best ligands in this subset. The numbers under the bars indicate how many studies contribute to the given value. For example, ethyl butyrate or 1-hexanol were covered in 9 studies, whereas ethyl hexanoate or benzaldehyde were only measured in 4 studies. Gray bars give the consensus values. White box plots right to the gray bars give median, quartiles (where available), and outliers (oval circles) obtained by using a leave-one-out strategy.

Note that the pattern changes somewhat for *Or59a* and changes dramatically for *Or59c* and *Or65a*. The most likely explanation is that for these receptors, the best ligands have not yet been found. Studies including more odors might find a better ligand, and targeted studies that exploit the combinatorial knowledge from the entire database might help. Nevertheless, it might also be that some receptors never reach the same strong responses as other receptors. In such cases, even though the individual best ligand has been found, the elicited response might still be weak as compared with maximal responses in other cells. With the globally scaled responses, it was possible to create response breadth plots for each single odor (Figure 3c), similar to the tuning breadth plots shown above. 2-Heptanone elicited responses in many receptors, some of which were negative. Methyl salicylate in contrast showed a very sharp profile evoking strong responses only in a few receptors.

Scaling odor responses across receptors is also a prerequisite for the creation of spatial odor response maps. In the *Drosophila* olfactory system, axons of sensory cells that express a given receptor converge stereotypically onto 1 glomerulus of the antennal lobe (AL), and thus, an activity map across receptor cells results in an activity map across olfactory glomeruli. These maps can be recorded directly, for example, using calcium imaging (Fiala et al. 2002; Wang et al. 2003; Silbering and Galizia 2007; Silbering et al. 2008). With the database presented here, virtual spatial activity maps in the antennal lobe can be generated; the map for 2-heptanone is shown in Figure 3d. On the webpage, the map for any of the odors in the database can be downloaded. The map visualizes activated glomeruli in shades of red, inhibited glomeruli in shades of blue, and indifferent glomeruli in white. Some glomeruli correspond to receptors, for which there is no response data; yet, in the case of 2-heptanone, these are the glomeruli D, DA1, and DC3 (see Figure 3d, light gray glomeruli). Other glomeruli do not have a value because the morphological mapping of these glomeruli onto a receptor is as yet unclear (e.g., glomerulus DP1m). Thus, the graphical display of these functional antennal lobes can also be used to earmark gaps in our knowledge of the *D. melanogaster* olfactome, gaps that need to be filled by targeted measurements. Interactive 3D renderings of these AL maps are also available from the Web site. A ball plot of OR response profiles is shown in Figure 3e for a subset (see also Supplementary Figure S8). Note that many entries are still missing, that is, unknown.

**Matching neurons, receptors, and glomeruli**

Odor response profiles in *D. melanogaster* have been measured in several ways: sensory cells that were identified morphologically, without knowing what receptor they expressed, expression of ORs in other receptor cells or heterologously, expression of calcium sensors in the receptor cells, and measurement of odor responses either in the dendrites or in the axon terminals. This diversity is possible because of a basic mapping property in this system: 1 receptor, 1 class of receptor cells, and 1 glomerulus. There are some exceptions to this scheme: some cells express more than 1 receptor, and some of the glomerular mapping strategies are more complex. Therefore, we included these cases into the database. The simplest one is given by *dOr22a*, which is coexpressed with *dOr22b*: because no function for *dOr22b* is known, only *dOr22a* has been mapped to the neuron ab3A and the glomerulus DM2. In cases where 2 receptors are coexpressed and each contributes to the odor response profile, we created a separate mapping for ORs (ligand-binding properties) and for receptor cells (odor response properties). For example, *dOr85e* and *dOr33c* are coexpressed in the receptor neuron pb2A (Goldman et al. 2005). The database contains 3 entries, but only the entry for pb2A is matched with glomerulus VC1 in the visualization of the antennal lobe. In this case, the functional relevance is high because the 3 odor response profiles differ.

**Mapping unlabeled response profiles into database**

In some cases, the mapping of receptor cell and receptor is not yet known. Here, the database can be used to find an appropriate match. To test this procedure, we used the database to find the receptor cell that expresses *dOr13a*. We expressed the calcium indicator G-CaMP under the control of the *dOr13a* promoter (Figure 4a,b) and recorded calcium odor responses to a total of 111 odors at a dilution of 1:100 (selected responses in Figure 4c, full results in Supplementary Table S3). For all odors that elicited responses, we further decreased the dilution in decadic steps until no responses were left. The best ligand was 1-octen-3-ol, and furfural elicited a calcium decrease (Nissler 2007). At this stage, the odor response profile of *dOr13a* was known, but the corresponding receptor cell was not. We thus used the consensus database to calculate how well the recorded response profile matched each of the known consensus response profiles. Data set ab6A had the best match (Figure 4d), which is a receptor neuron that had been characterized previously (de Bruyne et al. 2001) but for which the expressed receptor was not yet known. We also used a recently published data set in which odor responses in *dOr13a* were recorded (Kreher et al. 2008) and confirmed our result (data not shown). To confirm our link of *dOr13a* with ab6A, we mapped the area on the antenna where *dOr13a* is expressed (Figure 4a) and found that area to match the published location of ab6A (de Bruyne et al. 2001). The glomerulus that is innervated by neurons expressing *dOr13a* is DC2 (Couto et al. 2005) (Figure 4b). Thus, we conclude that ab6A expresses *dOr13a*, correcting previous suggestions that *dOr13a* might be expressed in intermediate sensilla (Couto et al. 2005). Taken together, we used a comparison between physiological recordings and the consensus database to find a match between receptor cells and receptor proteins and confirmed this by

**Figure 3** The complete consensus data set. **(a)** Tuning breadth plots (compare with Figure 1a) for 6 receptors based on the respective consensus data set. Note the pointed shape and negative responses in *Or59b* and *Or65a* and the broader profiles in *Or67a* and *Or67b*. Only few odor responses are available for *Or59c*. *n* Gives the number of odors but not the number of studies merged. Each receptor has been calculated separately and was therefore scaled independently of the other receptors. **(b)** Same as (a) but normalized across receptors (see text). *Or59a*, *Or59c*, and *Or65a* do not reach strong responses, indicating that these receptors have a different physiology or that the best ligands have not yet been identified. See Supplementary Figure S6 for additional plots. **(c)** Response breadth plots for 6 odors, that is, plotting responses against Or. Note that odors differ in their response breadth, for example, broad range

neuroanatomical analysis. A similar procedure might also be useful for interspecific studies, finding functionally homologous receptors across species.

### Estimating unknown receptor responses

As shown above, even with this comprehensive meta-analysis, our current knowledge of the *D. melanogaster* olfactome is quite incomplete. Thus, the database might lead to targeted studies toward a more complete olfactome. However, in several instances, it would be useful to have an estimate for an odor response even if none has been measured yet. Could the DoOR database be used for this purpose? We used local least squares imputation (Kim et al. 2005), which is a method for estimating missing values in a matrix (Supplementary Figure S4). As an example, Supplementary Figure S5a shows estimated responses in red. However, validating this approach using the leave-one-out technique, we found that this imputation technique is only reliable for a subset of odor responses (Supplementary Figure S5b,c; Wilcoxon test, $P = 0.5616$). Future studies will need to develop more appropriate algorithms for response estimation, possibly including external information such as chemical odor similarity.

### Relating olfactory space with other data

The *D. melanogaster* olfactome as it will be available with increasingly complete versions of the DoOR database can be used to answer several important questions in olfactory coding. As a teaser, we mention 4.

(1) Odor response properties can be mapped onto chemical space (Schmuker and Schneider 2007). In this approach, large data sets of chemical descriptors are used for characterizing chemicals, and multivariate statistics is used to extract those chemical descriptors that have the highest predictive values for odor responses of individual receptors or receptor families. This approach yields 2 very important results: first, it can be used to predict better ligands and/or unknown ligands for particular receptors (see above). Second, knowing which chemical properties best predict a receptor odor response profile can be used to understand mechanisms of ligand receptor interactions.

(2) Bioinformatic analysis of OR sequences. Mathematically, we have a similar approach as before, in which 2 related but distinct multidimensional spaces are compared and analyzed with respect to which parameters/

factors are most predictive for the interaction of the 2 spaces. Specifically, such a comparison might yield which sequence positions of the genes are correlated with odor response properties and which are not, thus generating hypotheses for odor-binding sites. Similar approaches have been taken for individual receptors, for example, the mouse *MOR42* subfamily and could be tested experimentally (Abaffy et al. 2007).

(3) Odor response properties can be mapped onto the behavioral meaning of odors: repellent or attractive odors (Semmelhack and Wang 2009) or pheromones and non-pheromones. Using the spatial representation of odor response patterns in the antennal lobe that can be generated from the DoOR package, it is possible to answer questions as whether behaviorally relevant odor responses are clustered and/or concentrated in particular antennal lobe areas or whether they are distributed and compare these results with experimental data.

(4) The logic of spatial arrangement of odor response properties in the antennal lobe can be analyzed. Supplementary Figure S3a shows an odor response similarity matrix for all glomeruli in the antennal lobe: some glomeruli have very similar odor response profiles (shown with dark red squares) and others are anticorrelated (blue). Is there a relationship between the spatial distance of glomeruli in the antennal lobe (Laissue et al. 1999) and their physiological similarity? We found the relationship to be significant, with a tendency of similar glomeruli to be closer neighbors (Supplementary Figure S3b), except when only cases with small odor counts (6:31) are considered. However, the slope of this relationship is small, accounting for 0.28 correlation value difference across the entire antennal lobe. We conclude that functional odor response properties have only a limited influence on the spatial location of glomeruli in the *D. melanogaster* antennal lobe, a conclusion that has significant implications for models of interglomerular computations in the antennal lobe (Galizia and Menzel 2001).

## Discussion

### The use of a functional atlas

Here, we create a functional atlas of odor responses for olfactory receptors, receptor cells, and olfactory glomeruli of

for 2-heptanone and isopentyl acetate and narrow range to methyl salicylate. *n* Gives the number of receptors included. **(d)** Physiological antennal lobe response to the odor 2-heptanone. By mapping each receptor to the glomerulus it innervates, we generate a fictive spatial response pattern in the antennal lobe. Excitatory responses are given in red and inhibitory responses in blue in 4 consecutive slices through the antennal lobe. UM, unmapped glomeruli, where the respective receptor is not yet known; NA, nonavailable glomeruli, where no odor responses have been measured for the corresponding receptor; BG, background material used for the shape of glomeruli beneath the indicated plane; D, dorsal; V, ventral; M, medial; L, lateral. Antennal lobe figure modified from Vosshall and Stocker (2007). **(e)** Plot of normalized odor responses across all available receptors, for a set of odors, including odors often used in behavioral studies in *Drosophila melanogaster*. Negative responses are given as empty circles. The complete table is in Supplementary Figure S8.
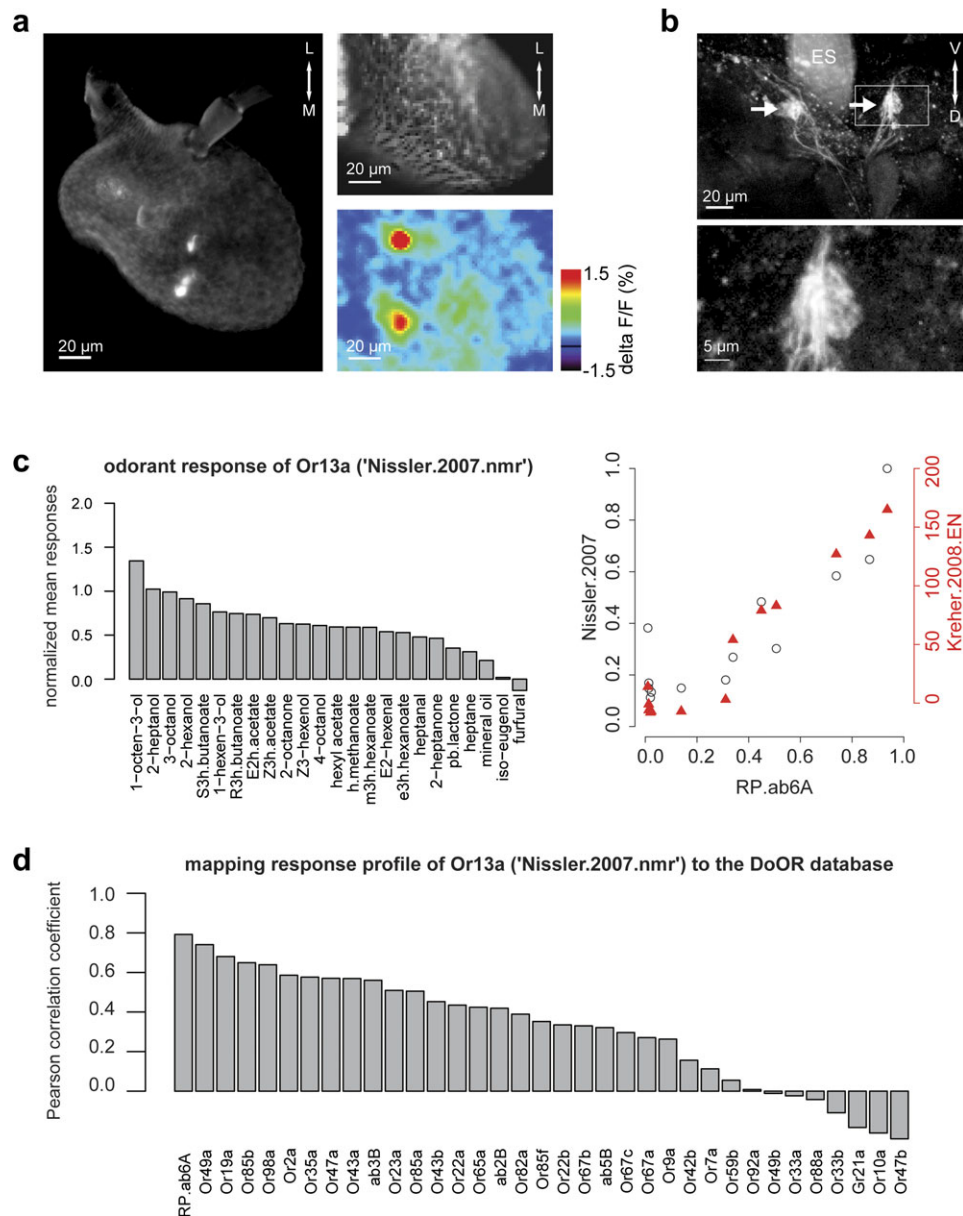
**Figure 4** Mapping response profiles to ORs. **(a)** Left panel: Confocal picture of the antenna of *Or13a*:GAL4;UAS:G-CaMP shows expression in a small number of olfactory sensilla. The location corresponds to that published for ab6A sensilla. Right upper panel: Anatomical picture of the antenna as seen in wide-field microscopy for calcium imaging. Right lower panel: False color–coded spatial response pattern to 3-octanol shows focalized responses. **(b)** Confocal picture of the antennal lobes of a *Or13a*:GAL4;UAS:G-CaMP fly shows fluorescence in 1 glomerulus for each antennal lobe (arrows), indicating that this Gal4 line targets a single receptor neuron population. The lower panel shows a magnification of the boxed area in the upper panel. ES, esophagus; D, dorsal; V, ventral. **(c)** Left: 24 selected odors that evoked calcium responses. S3h.butanoate, (*S*)-(+)-3-hydroxybutanoate; R3h.butanoate, (*R*)-(−)-3-hydroxybutanoate; E2h.acetate, E2-hexenyl acetate; Z3h.acetate, Z3-hexenyl acetate; h.methanoate, hexyl methanoate; m3h.hexanoate, methyl 3-hydroxyhexanoate; e3h.hexanoate, ethyl 3-hydroxyhexanoate; pb.lactone, gamma-propyl-gamma-butyrolactone; right: plotting the data measured knowing the receptor gene by calcium imaging (left ordinate) and electrophysiological recording (right ordinate) against the response data measured from ab6A (abscissa). **(d)** Pearson's correlation of the response profile over 111 odors measured by calcium imaging in *Or13a*:GAL4;UAS:G-CaMP flies to each known model response of antennal receptors. The best match was found with RP.ab6A.

the fruit fly *D. melanogaster*. This functional atlas represents a consensus data set combining all available data. It will serve as a reference work for olfactory physiologists, but it also represents a new approach of how to map different data sets onto each other. The only strict assumption made is that of a monotonic odor response function.

Most odors elicit a combinatorial pattern of activity across olfactory receptors, resulting in a stereotypical combinatorial pattern of activated glomeruli in the primary olfactory center (the mammalian bulb or the insect antennal lobe) (Galizia and Menzel 2001). In such a combinatorial system, the effect of removing individual receptors is difficult to

predict. For example, silencing *dOr22a* in *D. melanogaster* did not lead to a behavioral deficit in the response to any of the better ligands of this receptor, but it did create a deficit in response to a weak ligand (Keller and Vosshall 2007). This example shows that it is not sufficient to know the response of a single receptor class. Hence, the goal of this functional atlas is to generate the full olfactome of a species, in this case *D. melanogaster*. The data currently available do not yet include all receptors (see Supplementary Table S1), but the framework is open to new additions and will grow as more data will be collected by different laboratories.

Based on the complete olfactome, it will be possible to understand and to model the combinatorial nature of olfactory coding. In particular, the biological "olfactory space" can be derived from the data, that is, a description of how similar and dissimilar different odors are at the level of primary receptor input. At a later stage, when a similar database will be created for other species, it will be possible to compare these olfactory spaces for different species and thus to understand for what odors individual species have evolved higher resolution either in terms of discrimination capacity or in terms of sensitivity.

### The need of new mathematical tools

In principle, 2 approaches can be taken to create a complete functional atlas. In 1 approach, a mass screen using a dedicated technique would be used to create a homogeneous data set that results in a functional atlas. For example, in the visual system, the spectral response properties of photoreceptors can be mapped in great detail by electrophysiological recordings and once done the description is complete. Although attractive, this approach is not feasible in the olfactory system where the number of receptors is high in all species (*D. melanogaster* being among the most tractable) and the number of odors is infinite: every single study will always grasp but a partial view of the olfactome. Therefore, it is necessary to take the second approach, that is, to merge different data sets. Because these data sets differ in many respects, new mathematical tools are necessary. We have created a framework which allows for merging data sets of any kind as long as a single assumption is fulfilled: that the relationship be monotonic, that is, that better ligands in 1 study are expected to be better ligands in all studies (give or take variability).

This approach might also be useful in other studies where heterogeneous data sets need to be merged into metadatabases. Our entire package is open source. Without any change in the code, it can be adapted to the olfactory systems of other species: the only thing to do is to feed the data into a spreadsheet, create a graphical template for the antennal lobe output (if necessary), and a consensus database can be created. Thus, as soon as sufficient data will be available, the same platform will be usable to create olfactomes for other species, for example, mice or humans. With appropriate changes, the software could also be used for nonolfactory systems.

Although conceptionally and practically attractive, a database that is constantly evolving and including new data also creates problems: computational studies, for example, need to access standardized data sets because a change in the data set creates a situation where different results cannot be attributed unambiguously to a different model any more. Therefore, we will make older versions available indefinitely: the "DoOR 1.0" or "DoOR 2.0" will represent different stages in the publicly available data, such that computational studies will be able to consistently use a single reference olfactome, allowing for creating statistical or computational benchmarks.

### Limitations of the database

From a biological–physiological point of view, the data set presented here has 3 major drawbacks: it lacks information about 1) odor concentration, 2) complex stimuli, and 3) temporal response profiles. First, at the current stage, no information about responses to odor concentrations is included. This is a serious drawback because odor concentration is a fundamental parameter in olfaction. Some studies have measured odor responses across concentrations for all odors tested: in these cases, receptor responses can be coded as odor dilution that elicits half-maximal response strength (Pelz et al. 2006). In other studies, dose response curves were only measured for a subset of odors or not measured at all. For ligands with high affinity, this can create distortions in the database: for example, ethyl hexanoate and methyl hexanoate are currently the best-known ligands for *dOr22a* (Pelz et al. 2006). At high concentrations, however, the responses to these substances decrease due to fast receptor adaptation. Thus, in some studies that did not include dose response curves but tested many odors at high concentrations, these odors erroneously appear to be good, but not exceptional ligands. Some receptors have complex dose response curves for particular odors, further complicating the concentration aspect. Currently, there are not enough published data sets to include odor concentration into the database, but with an increasing number of studies, this will be possible. Including odor concentration as a parameter into the database will add 1 difficulty: measuring absolute odor concentration of a stimulus at the receptor cell in an experimental situation is not trivial. Thus, a concentration of 1:100 in 1 laboratory may not correspond to a concentration of 1:100 in another laboratory. Relative concentrations are less problematic: the relationship of 1:100 to 1:1000 will be 1:10 in all laboratories. Additional mathematical tools will be necessary to allow for automatic dose response curve shifts for data from different laboratories.

Second, complex stimuli are not covered in the database. These include odor mixtures but also other properties. For example, in a dynamical situation where odors are given as turbulent plumes, responses to some odors can be quite different as compared with the response to the same odor as

a single odor pulse (Schuckel et al. 2009). A related aspect needs to be considered for negative responses: many receptors respond to some odors with an activity decrease measured as a drop in firing rate or a decrease in intracellular calcium. However, some receptors have almost no spontaneous activity but might show inhibitory responses if activated beforehand. Here, an odor response is no longer a simple stimulus response property but rather dependent on previous activation. Such complexities cannot be covered in a functional atlas that is, in essence, a lookup table of simplified odor responses. However, these complexities are certainly important for the olfactory system and need to be considered in our quest to understand olfactory coding at large by generating dedicated physiological data sets.

Third, this functional data set maps odors to single values, disregarding the fact that odor responses are temporally structured at the level of olfactory receptors already. Response onsets to an odor have different time lags in different receptors, a property that could be included into the database as more data become available. Including more temporal information (e.g., phasic, phasic–tonic, tonic, or complex response patterns) will require additional tools. Temporal properties are more dependent on recording techniques than response magnitude: calcium imaging, intracellular recordings, or sensilla recordings might all reveal different aspects of the temporal complexity in a receptor neuron. Thus, including temporal information at the current stage would reduce the available data too much to make a consensus database useful.

Taken together, we present an open access software to assemble the complete olfactome of a species—here *D. melanogaster*. We hope that this service to the community will be of use for many further studies into olfaction of this and other species, and we will update the database as new odor response profiles become available.

## Supplementary material

Supplementary material can be found at http://www.chemse .oxfordjournals.org/. Additional material and the online version of DoOR is available at http://neuro.uni-konstanz.de /DoOR.

## Funding

## Acknowledgements

## References

Abaffy T, Malhotra A, Luetje CW. 2007. The molecular basis for ligand specificity in a mouse olfactory receptor: a network of functionally important residues. J Biol Chem. 282:1216–1224.

Benton R, Vannice KS, Gomez-Diaz C, Vosshall LB. 2009. Variant ionotropic glutamate receptors as chemosensory receptors in Drosophila. Cell. 136:149–162.

Buck L, Axel R. 1991. A novel multigene family may encode odorant receptors: a molecular basis for odor recognition. Cell. 65:175–187.

Couto A, Alenius M, Dickson BJ. 2005. Molecular, anatomical, and functional organization of the Drosophila olfactory system. Curr Biol. 15:1535–1547.

Crasto C, Marenco L, Miller P, Shepherd G. 2002. Olfactory receptor database: a metadata-driven automated population from sources of gene and protein sequences. Nucleic Acids Res. 30:354–360.

de Bruyne M, Clyne PJ, Carlson JR. 1999. Odor coding in a model olfactory organ: the Drosophila maxillary palp. J Neurosci. 19:4520–4532.

de Bruyne M, Foster K, Carlson JR. 2001. Odor coding in the Drosophila antenna. Neuron. 30:537–552.

Dobritsa AA, van der Goes van Naters W, Warr CG, Steinbrecht RA, Carlson JR. 2003. Integrating the molecular and cellular basis of odor coding in the Drosophila antenna. Neuron. 37:827–841.

Fiala A, Spall T, Diegelmann S, Eisermann B, Sachse S, Devaud JM, Buchner E, Galizia CG. 2002. Genetically expressed cameleon in Drosophila melanogaster is used to visualize olfactory information in projection neurons. Curr Biol. 12:1877–1884.

Galan RF, Weidert M, Menzel R, Herz AV, Galizia CG. 2006. Sensory memory for odors is encoded in spontaneous correlated activity between olfactory glomeruli. Neural Comput. 18:10–25.

Galizia CG, Menzel R. 2001. The role of glomeruli in the neural representation of odours: results from optical recording studies. J Insect Physiol. 47:115–130.

Glusman G, Yanai I, Rubin I, Lancet D. 2001. The complete human olfactory subgenome. Genome Res. 11:685–702.

Goldman AL, Van der Goes van Naters W, Lessing D, Warr CG, Carlson JR. 2005. Coexpression of two functional odor receptors in one neuron. Neuron. 45:661–666.

Hallem EA, Carlson JR. 2006. Coding of odors by a receptor repertoire. Cell. 125:143–160.

Hallem EA, Ho MG, Carlson JR. 2004. The molecular basis of odor coding in the Drosophila antenna. Cell. 117:965–979.

Herz AV, Meier R, Nawrot MP, Schiegel W, Zito T. 2008. G-Node: an integrated tool-sharing platform to support cellular and systems neurophysiology in the age of global neuroinformatics. Neural Netw. 21:1070–1075.

Keller A, Vosshall LB. 2007. Influence of odorant receptor repertoire on odor perception in humans and fruit flies. Proc Natl Acad Sci USA. 104:5614–5619.

Kim H, Golub GH, Park H. 2005. Missing value estimation for DNA microarray gene expression data: local least squares imputation. Bioinformatics. 21:187–198.

Kreher SA, Mathew D, Kim J, Carlson JR. 2008. Translation of sensory input into behavioral output via an olfactory system. Neuron. 59:110–124.

Kwon JY, Dahanukar A, Weiss LA, Carlson JR. 2007. The molecular basis of CO2 reception in Drosophila. Proc Natl Acad Sci USA. 104:3574–3578.

Laissue PP, Reiter C, Hiesinger PR, Halter S, Fischbach KF, Stocker RF. 1999. Three-dimensional reconstruction of the antennal lobe in Drosophila melanogaster. J Comp Neurol. 405:543–552.

Larsson MC, Domingos AI, Jones WD, Chiappe ME, Amrein H, Vosshall LB. 2004. Or83b encodes a broadly expressed odorant receptor essential for Drosophila olfaction. Neuron. 43:703–714.

Nissler A. 2007. Ligand search for genetically identified Drosophila olfactory receptors using calcium-imaging [thesis]. Konstanz (Germany): University of Konstanz.

Pelz D. 2005. Functional characterization of Drosophila melanogaster olfactory receptor neurons [Biologie: PhD thesis]. Berlin (Germany)]: Freie Universität Berlin.

Pelz D, Roeske T, Syed Z, de Bruyne M, Galizia CG. 2006. The molecular receptive range of an olfactory receptor in vivo (Drosophila melanogaster Or22a). J Neurobiol. 66:1544–1563.

R Development Core Team. 2009. R: a language and environment for statistical computing. Vienna (Austria): R Foundation for Statistical Computing.

Schmuker M, de Bruyne M, Hähnel M, Schneider G. 2007. Predicting olfactory receptor neuron responses from odorant structure. Chem Cent J. 1:11.

Schmuker M, Schneider G. 2007. Processing and classification of chemical data inspired by insect olfaction. Proc Natl Acad Sci USA. 104: 20285–20289.

Schuckel J, Torkkeli PH, French AS. 2009. Two interacting olfactory transduction mechanisms have linked polarities and dynamics in Drosophila melanogaster antennal basiconic sensilla neurons. J Neurophysiol. 102:214–223.

Semmelhack JL, Wang JW. 2009. Select Drosophila glomeruli mediate innate olfactory attraction and aversion. Nature. 459:218–223.

Silbering AF, Galizia CG. 2007. Processing of odor mixtures in the Drosophila antennal lobe reveals both global inhibition and glomerulus-specific interactions. J Neurosci. 27:11966–11977.

Silbering AF, Okada R, Ito K, Galizia CG. 2008. Olfactory information processing in the Drosophila antennal lobe: anything goes? J Neurosci. 28:13075–13087.

Smart R, Kiely A, Beale M, Vargas E, Carraher C, Kralicek AV, Christie DL, Chen C, Newcomb RD, Warr CG. 2008. Drosophila odorant receptors are novel seven transmembrane domain proteins that can signal independently of heterotrimeric G proteins. Insect Biochem Mol Biol. 38:770–780.

Stensmyr MC, Giordano E, Balloi A, Angioy AM, Hansson BS. 2003. Novel natural ligands for Drosophila olfactory receptor neurones. J Exp Biol. 206:715–724.

Vosshall LB, Amrein H, Morozov PS, Rzhetsky A, Axel R. 1999. A spatial map of olfactory receptor expression in the Drosophila antenna. Cell. 96: 725–736.

Vosshall LB, Stocker RF. 2007. Molecular architecture of smell and taste in Drosophila. Annu Rev Neurosci. 30:505–533.

Wang JW, Wong AM, Flores J, Vosshall LB, Axel R. 2003. Two-photon calcium imaging reveals an odor-evoked map of activity in the fly brain. Cell. 112:271–282.