The Journal of Experimental Biology 213, 2461-2475

doi:10.1242/jeb.036079

A syntax of hoverfly flight prototypes

Bart R. H. Geurten[1,*], Roland Kern[1,2], Elke Braun[1,2] and Martin Egelhaaf[1,2]

[1]Neurobiology, Bielefeld University, PO Box 10 01 31, D-33501 Bielefeld, Germany and [2]Center of Excellence 'Cognitive Interaction Technology', Bielefeld University, PO Box 10 01 31, D-33501 Bielefeld, Germany

*Author for correspondence (bart.geurten@uni-bielefeld.de)

Accepted 23 March 2010

SUMMARY

Hoverflies such as *Eristalis tenax* Linnaeus are known for their distinctive flight style. They can hover on the same spot for several seconds and then burst into movement in apparently any possible direction. In order to determine a quantitative and structured description of complex flight manoeuvres, we searched for a set of repeatedly occurring prototypical movements (PMs) and a set of rules for their ordering. PMs were identified by applying clustering algorithms to the translational and rotational velocities of the body of *Eristalis* during free-flight sequences. This approach led to nine stable and reliable PMs, and thus provided a tremendous reduction in the complexity of behavioural description. This set of PMs together with the probabilities of transition between them constitute a syntactical description of flight behaviour. The PMs themselves can be roughly segregated into fast rotational turns (saccades) and a variety of distinct translational movements (intersaccadic intervals). We interpret this segregation as reflecting an active sensing strategy which facilitates the extraction of spatial information from retinal image displacements. Detailed analysis of saccades shows that they are performed around all rotational axes individually and in all possible combinations. We found the probability of occurrence of a given saccade type to depend on parameters such as the angle between the long body axis and the direction of flight.

Key words: hoverfly, flight behaviour, clustering, syntax, movement segregation, saccades.

INTRODUCTION

The movements of animals and humans are, in many cases, fast, complex and continuous. Often such behaviour is segregated into ordered components (Flash and Hochner, 2005; Jenkins and Mataric, 2003; Mussa-Ivaldi and Bizzi, 2000). For example, in the face-grooming behaviour of mice, distinct movements were identified and their ordering analysed (Fentress and Stilwell, 1973). In computer science movement templates are used to recognise human sign language (Liang and Ouhyoung, 1998) and to analyse the movements of players in complex simulation games (Thurau and Hlavac, 2007). Also in robotics or in developing computer games, simple movements are sequenced to facilitate the generation of naturalistic locomotion patterns. The common ground for these examples, which are taken from behavioural analysis, computer-based recognition and machine motion planning, is the segregation of complex and continuous behaviour into simple consecutive building blocks of movement.

In many species a segregation of flight sequences into reoccurring rotational and translational flight segments has been described (Bender and Dickinson, 2006; Boeddeker and Hemmi, 2010; Boeddeker et al., 2010; Collett and Land, 1975a; Eckmeier et al., 2008; van Hateren and Schilstra, 1999; Mronz and Lehmann, 2008; Schilstra and van Hateren, 1999; Wagner, 1986). This finding motivated us to investigate whether the flight manoeuvres of free-flying hoverflies (*Eristalis tenax*, Linnaeus) can be segregated into more detailed consecutive movement components.

We chose hoverflies for our analysis because of their rich repertoire of extremely fast and virtuosic flight manoeuvres, even under spatially constrained conditions. They can move in nearly every direction and their translational velocities range from 10 m s^{-1} to an almost total lack of movement in mid-air (hovering) (Collett and Land 1975a). Moreover, its rather small brain with less than a million neurons makes *Eristalis* an interesting model organism for subsequent neurophysiological experiments in motion vision (Barnett et al., 2007; Nordström et al., 2008; O'Carroll et al., 1996; O'Carroll et al., 1997).

We used clustering algorithms to identify movement components within flight trajectories of *Eristalis*, which we call prototypical movements (PMs) in this article. Clustering allows us to segregate behaviour without having to predefine behavioural components. Instead, this approach relies on the assumption (1) that the behavioural data can be represented quantitatively in a suitable way, and (2) that an appropriate distance measure can be found that attributes small distances to similar behaviours. Merely identifying PMs is not sufficient for understanding behavioural control. In the context of visually guided flight behaviour of *Eristalis*, one might also wish to know the rules for the arrangement of PMs into complex flight trajectories. By calculating transition probabilities between PMs, such rules may be derived, and can be regarded as the grammar of a formal language (Chomsky and Schützenberger, 1963) in which the PMs are treated like an alphabet. In combination, this grammar and alphabet give us a syntax of movement.

How variable is this syntax in the real world? In monkeys, prototypical limb movements elicited by prolonged microstimulation of the motor cortex are not constant, but vary depending on the starting positions of the joints (Graziano et al., 2002; Graziano et al., 2005; Stepniewska et al., 2005). This finding motivated our investigation of whether PMs of *Eristalis* also vary with the situational context. We analysed whether this is the case for PMs that represent fast rotational movements called saccades. Saccades are sometimes described or modelled as a fixed motion pattern, as for example in the fruit fly Drosophila (Mronz and Lehmann, 2008; Tammero and Dickinson, 2002) or as a flexible

motion pattern with varying top speed and amplitude, as in other flies (Collett and Land, 1975a; Schilstra and van Hateren, 1999; Lindemann et al., 2008). Therefore, we analysed saccades to discover whether the PMs of *Eristalis* are influenced by situational parameters and if they show the same flexibility as the aforementioned movements of monkeys.

MATERIALS AND METHODS

All calculations presented in this article were done with MatLab R2008b (The Mathworks Inc., Natick, MA, USA).

Data acquisition

We observed freely flying hoverflies *Eristalis tenax* Linnaeus collected close to Bielefeld University. The animals were filmed with two black-and-white high-speed cameras (Motion Pro, Redlake, Eningen, Germany) at 500 frames s^{-1}. The cameras were orthogonally orientated, so that one camera filmed the arena from above and one from the side. The camera positions were calibrated with Jean-Yves Bouguet's MatLab camera calibration toolbox (www.vision.caltech.edu/bouguetj/calib_doc/) (Bouguet, 1998) to derive the 3D position of the animal from the two 2D positions within the images. We used two differently sized flight arenas, a small cubic arena (142×242×172 mm^3) and a large cylindrical arena (400 mm diameter, 700 mm height). Both arenas were wallpapered with red and white random cloud patterns with a 1/*f* naturalistic spatial frequency distribution (van der Schaaf and van Hateren, 1996). The colour red was chosen because it provides a good contrast against white for the fly, and good contrast between fly and background for the camera, which simplifies image processing. In the small arena we used a Perspex ceiling and front wall to obtain footage from the complete arena. In the large arena we cut a hole (20 cm diameter) in the wallpaper at the side and used a Perspex ceiling, allowing us to film the flies. We recorded 86,983 frames (*ca*. 3 min) in the small arena from nine flies. The 47 recorded flight sequences had a mean length and standard deviation of 3.63±0.59 s. The 33 flight sequences recorded in the large arena sum to 50,622 frames (*ca*. 1 min 40 s) and have a mean length of 3.01±1.01 s. In the large arena we used eight flies. Different flies were used for the two arenas.

Determination of position and orientation in space

The fly's body position as well as its yaw and pitch orientation in space were calculated by tracking characteristic points of the fly in each pair of corresponding 2D images and determining the 3D position using the camera calibration information. The chosen tracking points were the front of the head and the tip of the abdomen (see Fig. 1A). Tracking was done using the custom-made software package FlyTrace (E.B. and J. P. Lindemann, www.uni-bielefeld.de/biologie/Neurobiology/). We filtered the 2D trajectories for each camera with a Butterworth filter (2nd degree; relative cut-off frequency 0.1), to reduce jitter introduced by auto-detection. We tested different filter settings by overlaying filtered trajectories with the original footage until we found a filter setting that both eliminated jitter from the detection process and matched the fly's position. From the eventually derived pair of corresponding 2D trajectories, the 3D trajectory was calculated using the stereo-triangulation method in Jean-Yves Bouguet's camera calibration toolbox (Bouguet, 1998). Treating the front of the head and the tip of the abdomen independently allowed us to determine the body position and the yaw and pitch angles of the body long axis in space. This method does not deliver the roll angle of the fly.

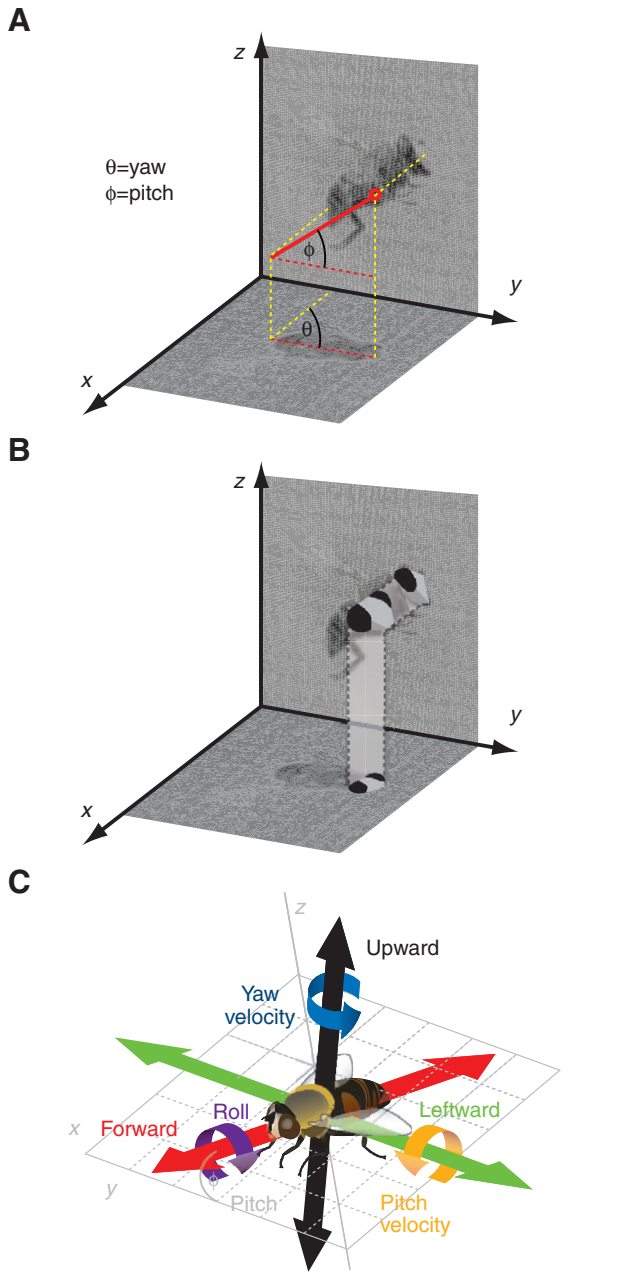

Fig. 1. Image analysis and coordinate system. (A) Method used to derive the body position from the video frames of two synchronised, orthogonally oriented video cameras. Single characteristic points on the head and abdomen, respectively, were used. By triangulation of the two points the orientation and position of the body length axis was reconstructed. Note that the roll angle cannot be calculated with this method. (B) Acquiring head position. A 3D model of the head was adjusted to match the appearance of its projection in corresponding video frames. The model was adjusted in position and orientation for all six degrees of freedom. (C) Definition of velocities used for analysis. Velocity vectors are fly centred. The velocities are shown in the colours we used throughout this article for the prototypical movements (PMs).

There is always a trade-off between the size of the observation area and spatial resolution. Therefore, head tracking was only possible in the small flight arena. In a limited number of frames, we manually fitted a 3D model to the fly's head, using FlyTrace,

in order to track head positions and orientation (see Fig. 1B). This 3D model method allowed us to derive all six degrees of freedom.

Describing behaviour by velocities

We assumed the translational and rotational velocities of the fly in a fly-based coordinate system constitute characteristic parameters of flight behaviour. Following conventions in computer science, we call these parameters 'features'. The forward velocity is defined as the velocity along the body long axis. The sideways velocity is orthogonal to this vector in the horizontal plane. Orthogonal to the plane formed by the two preceding velocity vectors is the upward velocity vector. In addition, the rotational velocities around these three axes were taken into account (see Fig. 1C). The velocities were calculated from the difference in position and orientation between two frames. The roll velocity could only be extracted for the small set with available head data. We chose these velocities because they are widely used to characterise insect flight. Although other features could be used, we wanted to keep our work comparable to other studies of the field. We also conducted an analysis in a world-based translational velocity system. The results were very similar and there was no qualitative difference between conclusions for the two different velocity sets.

In some parts of the analysis, the ground speed was defined as the projection of the 3D translation velocity vector into the horizontal plane. Neither roll velocity nor ground speed was used in the clustering analysis described below.

Finding prototypical movements by cluster analysis

To characterise flight behaviour, we took the set of rotational and translational velocity values corresponding to each point of the trajectory. Repeatedly occurring similar movements lead to similar velocity data. To quantify similarities we decided to use the squared Euclidean distance as a measure. The Euclidean distance is in common use for calculating distances between continuous real values, as our velocity values are. The squared Euclidean distance delivers the same results when comparing distances between data points with each other, as we did here, and is computationally simple. With this similarity measure, the velocity feature data points representing different reoccurring movements constitute dense clouds of data points within the high-dimensional feature space that are distinct from each other. For analysing these structures we used cluster analysis. This approach is able to robustly distinguish the clouds, also called clusters. For doing so, the feature data space is segregated into classes by minimising the distances between the data points within the classes and maximising distances between classes. Each class can be characterised by its cluster centroid, which is the point in data space that has minimal distance to all other data points in that cloud. If the feature data space provides significant structures due to reoccurring movements described by selected velocity values, this procedure provides a stable set of classes and an assignment of each velocity data point to one class.

Assigning the corresponding class label to the individual data points within the temporal sequence of the flight velocity data leads to a sequence of class labels that can be easily segregated by identifying subsequences of constant class labels. Such a subsequence of constant class labels '*x*' represents what we call a PM, e.g. subsequence '*xxxxx*' is collapsed into a single PM with an extended duration. The length of each subsequence determines the duration of the corresponding PM.

Whether a set of classes constitutes an appropriate representation of the feature data and therefore delivers meaningful movement components has to be evaluated in a post-processing procedure. In the following sections the details of preparing feature values for clustering, the clustering procedure itself and the evaluation step are briefly described (for details, see Braun et al., 2010).

Normalisation

For calculating similarities between feature values as required for clustering, the feature values have to be normalised, because translational and rotational velocities are given in different physical dimensions. Moreover, because the individual translational and rotational velocity components greatly differ numerically (e.g. yaw velocities can be much larger than pitch velocities), all data of a given feature were grouped and normalised independently. In this way each velocity dimension contributes equally to the clustering independent of its relative variability. All velocity dimensions get equal importance for the clustering, which generally provides different clustering results in comparison to clustering without normalisation.

We used a standard normalisation technique called *z*-score which normalises the data to zero mean and a standard deviation of one, and leads to non-dimensional variables that are suited to clustering. We chose this normalisation technique because most of the data are concentrated in the numerical range –1 to 1 and outliers are not overemphasised as occurs for example by normalisation to the maximum.

Agglomerative hierarchical pre-clustering

With this approach we determined candidates for the suitable number of clusters k to be identified within the feature data. The approach starts by treating every feature data point as an individual cluster, iteratively merges two clusters into a new one that minimises a given criterion and stops when all data are contained within one cluster. For determining an appropriate number of clusters Ward's criterion is used. This criterion treats the increase in inner cluster variance resulting from the merging of two clusters as merging costs. Identifying steps in the sequence of merging costs that is assigned to each intermediate number of clusters indicates candidates for the suitable number of clusters k. The hierarchical clustering approach evaluates locally pairs of data points and clusters which makes it applicable only to small data sets. We applied hierarchical clustering to parts of our large data set (10% of the original data set) in order to determine suitable numbers of clusters to be tested with the entire data set by the *k*-means method.

k-means clustering

With the *k*-means clustering algorithm we partitioned our large sets of noisy high-dimensional feature values into k clusters. This approach has proven to be appropriate because of its robustness and computational simplicity. The clustering approach needs the number of clusters k to be a pre-defined parameter (MacQueen, 1967). To define a value k that leads to a meaningful segregation of the data, we evaluated the results of the agglomerative hierarchical clustering described above.

For a given number of clusters k, *k*-means partitioning aims at minimising the overall sum of distances between all the data points and their respective assigned centroids. The calculation of this partitioning is done iteratively. After choosing starting positions for the centroids, all feature data points that are closer to centroid $x_i \in x_{i,k}$ than to any other centroid become members of the cluster of centroid x_i. Based on these intermediate clusters, the positions with minimal distance to the feature data points are calculated individually for each cluster and they become the new centroids. Since all centroids have been moved, the cluster members are

recalculated on the basis of the current centroid positions. This is done until the centroids no longer move significantly or another user-defined criterion (e.g. maximal number of iterations) is satisfied. Note that the *k*-means algorithm may get caught in different local minima depending on the starting positions for the centroids. Without any knowledge about suitable starting positions, they are selected randomly from the feature data set. We stabilise these random positions by *k*-means pre-clustering for a subset of 10% of the data and use the centroids resulting from the pre-clustering as new starting positions. In order to further compensate for the dependence of the results on the randomised starting positions, we calculated 10 runs and chose results where the centroids provide minimal distance to their assigned feature data points. We used the MatLab R2008b *k*-means implementation with squared Euclidean distances.

To use the *k*-means approach for classifying feature values we have to determine the most appropriate number of clusters *k*. This implies we need to determine which clusters represent the feature data best. By evaluating several numbers of clusters, as described in the following paragraph, we are able to select a model and, furthermore, estimate whether this model is appropriate and leads to meaningful prototypical movements.

Evaluation

For determining the number of clusters that best segregate our data and for evaluating whether they constitute an appropriate representation, we evaluate the clustering result by two criteria. The quality of clustering assesses whether the clusters represent significantly distinguishable clouds of feature data points. The instability reflects the reproducibility of the clusters with respect to varying starting conditions and changes of the data set. Both criteria are chosen to evaluate general characteristics of an appropriate clustering instead of evaluating the different classification results because we do not have prior knowledge of the 'classification truth' to compare with.

The variable conditions to be evaluated were achieved by randomly choosing starting positions of the centroids and by leaving out continuous chunks of the data (10%, 25% and 50%). We also performed an evaluation with the complete data set, varying randomly chosen starting positions. Through optimising quality and stability we find the number of clusters that best segregate our data. Every number of clusters between 2 and 20 was tested. For computational reasons, for cluster numbers between 20 and 50 only every fifth number was tested.

For each chunk size of omitted data we compared different runs for a given number of clusters. For this comparison, each two resulting cluster responses were matched by a minimal cost-perfect matching algorithm (Kuhn, 1955; Kuhn, 1956). The cost of this matching gives us a metric for the dissimilarity of the two clustering results. We call this measure instability.

Furthermore, we analysed the quality of clustering by analysing whether the data clouds are dense and distinguishable from each other. Hence we determined for each cluster the similarity of its assigned data points in relation to the distances of the cluster centroid to the others. Our quality measure is the shortest squared Euclidean distance between the cluster centre and any neighbouring cluster centre (outer distance) divided by the mean squared Euclidean distance between all data points of the cluster (inner distance). By taking the mean inner distance we get a value for the variance inside the cluster. A large outer distance in relation to a small inner distance indicates distinct behaviour and leads to high quality values. We calculated the quality of a clustering result by taking the mean quality of its individual clusters. For determining the quality for the different runs within one leave-out size condition (continuous chunks) and a given number of clusters we also calculated the mean value.

Evaluating the clustering results with respect to their stability and quality for a varying number of clusters allowed us to identify the number of clusters that best segregate the data. The optimal number of clusters is found by searching for a combination of low instability and high quality. Identifying stable clusters of sufficient quality is a precondition for clusters to represent significant structures within the feature data.

After we determined the best number of clusters, we had to select the set of *k* clusters out of the sets resulting from the different runs to be further analysed. We took the set with the lowest matching cost (when compared with all other cluster sets).

Starplots

The cluster centroids of our data are five-dimensional vectors corresponding to five features. They are visualised as star plots. Every feature dimension provides its own coloured axis within the unit circle.

As a consequence of normalisation, the centroids are non-dimensional values. To interpret their components we have to de-normalise them to restore the physical values. In order to visualise translational and rotational velocities within a given star plot, we normalise the three translational velocities to the combined maximum of their absolute values. This was done for the two rotational velocities in an equivalent way. Note that yaw velocities are much faster than pitch velocities and by this normalisation, pitch velocities and especially differences between different pitch velocities might look very small.

Markov analysis

A level one Markov analysis consists of a set of states and the transition probabilities between them. The set of states is given by the set of PMs. Transition probabilities can be derived from the sequence of PMs. We statistically tested the transitions against the null hypothesis of chance levels of transition probabilities between PMs. This procedure takes into account the fact that, if a PM occurs very often, the probability of transition to this PM from any other PM is expected to be high even if the transitions occur by chance. We calculated a confidence interval for each transition probability using Bernoulli statistics. If the global distribution value of the target PM was inside the confidence interval we called it a chance level transition. Transitions significantly below or above this chance level were treated as 'prohibited' or 'permissible' transitions, respectively ($P<0.05$).

Temporal detail analysis of velocity traces during saccades

The velocity threshold for saccades was set to $400\,\mathrm{deg\,s^{-1}}$ for yaw, $150\,\mathrm{deg\,s^{-1}}$ for pitch and $75\,\mathrm{deg\,s^{-1}}$ for roll. The peak velocity of a saccade was found by a maximum search in between the two crossings of the threshold. The start and end of a saccade were defined by the next inflection or crossing of zero velocity.

RESULTS

Protoypical movements in differently sized flight arenas

We filmed flights in two arenas, a small cubic arena and a much larger cylindrical arena. The cylinder had a volume 21 times bigger than the cubic arena. In each frame we analysed the position and orientation of the fly (Fig. 1A) and calculated the velocities within the fly-based coordinate system. The ground speed in the large arena

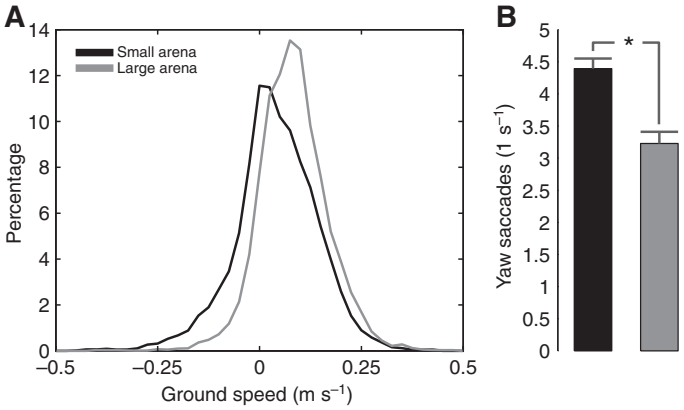

Fig. 2. Ground velocity distribution and saccade frequency. (A) Ground velocity distribution as a percentage for the large and small arena (bin size 0.01 m s^{-1}). Negative speeds correspond to backward flight. Data were tested with the Kruskal–Wallis test and were significantly different. (B) Mean (±s.d.) saccade frequency for the large (grey) and the small (black) arena. Means differ significantly from each other (*P<0.05). The large arena is 21 times larger in volume than the small arena.

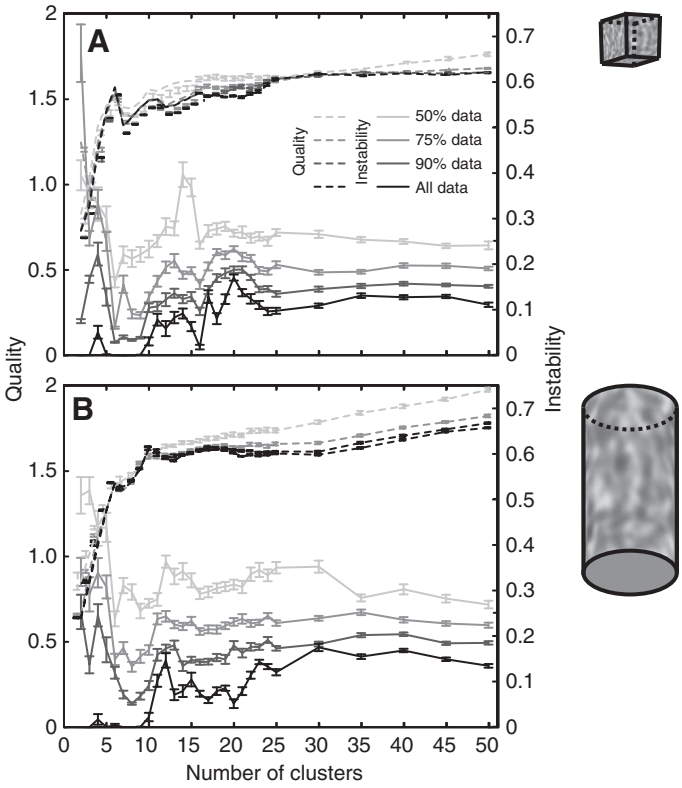

Fig. 3. Quality and instability of *k*-means clustering. (A) Small arena, (B) large arena. Overall quality (dashed lines) and instability (solid lines) are plotted against the number of clusters used. The overall quality value is the mean quality value of all cluster runs. The mean quality value of a cluster run is the mean quality of all clusters in a cluster run. The quality value for each cluster is the fraction between the smallest outer distance and the mean inner distance. Inner distance is defined as the mean distance between all points of a cluster. Outer distance is the smallest distance between the centroid and all the other centroids. The overall instability value is the mean instability of all cluster runs of the same condition. The mean instability of a cluster run is the mean instability of all clusters, i.e. the cluster set in that run. The instability of a cluster set is defined as the mean of the minimal cost perfect matching for this set of clusters to all other sets of cluster runs of the same cluster number and data size. Differently sized chunks of the data set (body trajectories) were used (see legend). If less than 100% of the data were used, the data were cut out in 50 different positions in a round robin fashion.

tended to be higher than in the small arena (Fig. 2A). In both arenas backward flight occurred frequently. Interestingly, *Eristalis* flew backwards more often in the small arena (25% of the total flight time) than in the large one (13% of the total). Also, saccades were exhibited more often in the small arena (Fig. 2B). It appears that the more confined space in the smaller arena forces the flies to approach walls more frequently, whereupon they may retreat, flying backwards and turning away from the walls. Note that both arenas allow the animals to show only a small part of their behavioural repertoire. For instance, velocities of up to 10 m s^{-1}, as were previously described in other settings (Collett and Land, 1978), were not observed in our flight arenas.

To describe the flight behaviour of *Eristalis* in a more general way, we searched for PMs by applying the clustering approach as described in Materials and methods. PMs were determined separately for the two arenas. The hierarchical pre-clustering yielded a range of between 2 and 50 clusters for both arenas. We tested these numbers of clusters with the *k*-means algorithm and evaluated the instability and quality of the clustering results (see Fig. 3). To this end, we left out 50 different chunks of data of different sizes (0%, 10%, 25% and 50%). We found that we could increase the number of clusters to up to nine in both arenas and still get stable results for the complete data sets. Leaving out bigger chunks of the data and comparing the results obtained for the different data sets led to more instability, but the mean set of centroids stayed the same irrespective of the amount of data left out. However, in all cases there was a local minimum at nine clusters. The quality reached a plateau at nine clusters for data sets of all sizes, suggesting that quality cannot be improved much by increasing the number of clusters (see Fig. 3). We could also have chosen six clusters, which have similar instability and quality values to those of nine clusters. Since we tried to derive the classification with the highest resolution that is possible under the constraints of quality and instability, we selected nine clusters as appropriate for both flight arenas for our further analysis.

k-means clustering treats the velocity values of a frame independent of previous and subsequent frames. The clustering process thereby omits the information about the time course of the flight. Classifying each frame of the sequence individually raises the possibility that PMs were derived from single frame events, which we have to treat as classification errors, given our high temporal resolution of 2 ms per frame. We calculated the mean duration of PMs *via* the mean number of frames individually assigned to them (see Fig. 4). The duration values were a first indication that the clustering was successful, on average, because, even though clustering omits the temporal structure of the data, the results show that structure again. Essentially the same types of nine PMs were found for the small and the large arena (see Fig. 4).

We can clearly distinguish PMs that contain no or only translational velocities and those that contain significant rotational components (PM1, 2, 8). In both arenas, PM1 and PM2 (for identification numbers see lower right corner of the centroid plots in Fig. 4) describe saccadic yaw turns that are combined with a reduction of the pitch angle and a slight drop in altitude as well as a sideways movement in the direction of the saccade. These two

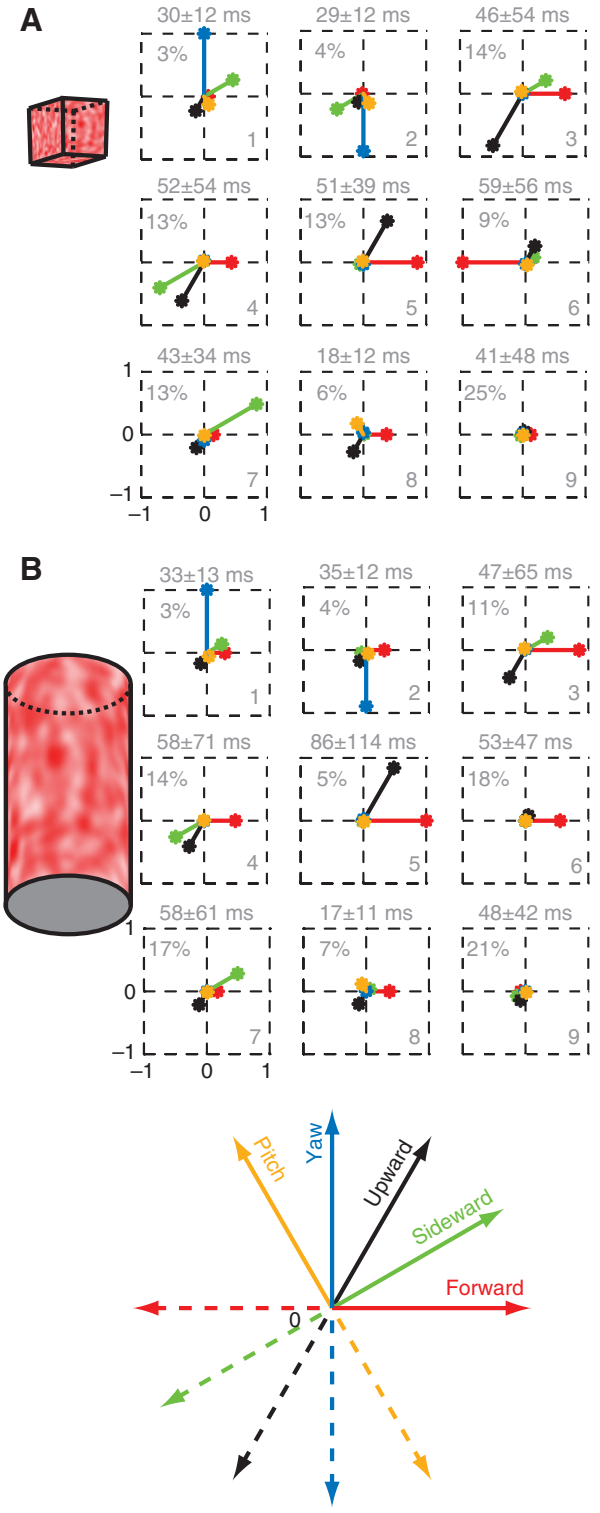

Fig. 4. PMs for body trajectories in both arenas. The centroids of the *k*-means clustering of fly body velocities are regarded as PMs and plotted as stars. Each velocity has its own colour-coded axis, as shown in the key at the bottom. For illustration of the velocities used see Fig. 1C. Dashed axis parts in the key mark negative values. The velocities were normalised as follows. All translational velocities were normalised to their joint absolute maximum. Rotational velocities were normalised analogously. In the lower right corner of each centroid plot the identification number used in the main text is given in grey. The mean (±s.d.) duration of a centroid within the trajectory data is given above each plot and the percentage of data assigned to the centroid is in the upper left corner of the plot. Centroids of flights in the small arena (A) and in the large arena (B) are shown. The arenas are shown at the side in relative size to each other.

PMs make up only a small fraction of the data (3–4%; see Fig. 4). They also have the smallest standard deviation, compared with their own duration, as found by analysing the mean durations of the PMs (Fig. 4).

PM3 and PM4 describe a combined forward and sideways movement. Note that movements in a horizontal plane are a combination of forward movement along the body long axis (red line in Fig. 4) and downward velocity (black line in Fig. 4), because the body was pitched upwards by at least 25 deg during all observed cruising flights.

PM5 corresponds to a combined upward and forward movement in which altitude is gained. This PM is more often exhibited in the small arena than in the large one, but, on the other hand, it has a longer duration, on average, in the large arena (see Fig. 4).

The only PM that differs qualitatively between the arenas is PM6. In the small arena it reflects backward motion. In contrast, in the large arena it describes a pure forward motion. The backward PM characteristic of the small arena is less often exhibited than the forward PM found in the large arena. Although flies do exhibit backward motion in the large arena too (see Fig. 2), these events are too rare and not sufficiently distinct to form a PM on their own. Instead another forward PM is formed for the large arena. If we increase the number of clusters from 9 to 12 in the *k*-means algorithm, we also find a backward cluster in the set of PMs for the large arena (not shown). Although the set of clusters is fairly unstable, if 12 PMs were taken into account, the backward PM was formed in every trial of the *k*-means analysis. Under the given constraints of nine clusters, the backward PM is not formed from the data derived for the large arena.

PM7 reflects motion mainly in a leftward direction. Surprisingly, this PM has no mirror PM like PM1/PM2 and PM3/PM4. Closer inspection of the sideways movement distribution between the PMs shows that rightward movement is split across the other prototypes. Hence whenever the fly is moving exclusively to the right, this movement will be grouped into one of the other prototypes, most often into PM4. In the small arena PM7 seems to represent those leftward velocities which are higher than the leftward movement component in PM3. PM7 and PM3 divide the leftward velocity into a fast and a slow class whilst PM4 represents the mean value of rightward velocity.

PM8 describes another form of saccade, a pitch saccade. It shows a rather definite duration, similar to PM1 and PM2. In the star plots the pitch velocity appears not to be much faster in PM8 when compared with the other PMs. This is because the rotational velocities were normalised against their overall maximum, which was set by the much larger yaw velocity of 1131 deg s^{-1} (small arena). Nonetheless, the pitch saccade is still very fast at 228 deg s^{-1} (small arena). There is no pitch PM mirroring PM8, as for PM7. Again, the downward movement is split up into several PMs, namely PM1 and PM2. Saccades are inspected in more detail below.

The final PM is hovering. PM9 makes up more than 20% of the flight time in both arenas (see Fig. 4). Its duration is highly flexible.

Even though key features such as ground velocity and saccade frequency change with arena size, PMs stay qualitatively the same,

Table 1. Quality and instability of the resulting nine clusters for the large and the small arena

	Small arena					Large arena				
	Quality			Stability		Quality			Stability	
PM	Inner	Outer	Fraction	Mean	s.d.	Inner	Outer	Fraction	Mean	s.d.
1	6.17	11.40	1.85	0.0002	0.0002	7.20	12.73	1.77	0.0002	0.0003
2	5.96	10.57	1.77	0.0002	0.0003	5.62	10.51	1.87	0.0000	0.0000
3	1.87	2.72	1.45	0.0010	0.0011	2.41	3.35	1.39	0.0004	0.0009
4	2.23	2.94	1.32	0.0006	0.0006	1.95	2.92	1.50	0.0003	0.0003
5	2.12	2.45	1.16	0.0005	0.0005	5.60	5.20	0.93	0.0002	0.0003
6	2.81	3.43	1.22	0.0004	0.0004	1.28	2.67	2.09	0.0014	0.0018
7	1.98	2.61	1.32	0.0008	0.0008	1.77	2.72	1.54	0.0005	0.0005
8	4.32	3.74	0.87	0.0021	0.0031	3.85	4.02	1.04	0.0011	0.0013
9	1.17	2.45	2.10	0.0006	0.0006	1.81	2.67	1.48	0.0004	0.0005

The inner distance is the mean squared Euclidean distance between all data points of the cluster. The outer distance is the smallest Euclidean distance between the centroid (the prototypical movement, PM) and all other centroids. Stability is the mean cost of all the minimal cost perfect matchings for that cluster to the other 49 cluster runs for the complete data set.

Table 2. Transition probabilities as a percentage of the prototypical movements from flights in the small arena

PM	1		2		3		4		5		6		7		8		9	
1	0	x	2	○	26	+	2	○	3	○	5		40	+	9	○	14	○
2	0	○	0	x	0	○	50	+	0		14	+	2	○	9	○	26	
3	6		3	○	0	x	14	+	6	○	0	○	28	+	31	+	14	○
4	3	○	11	+	23	+	0	x	6	○	3	○	0	○	23	+	31	+
5	3	○	8	+	3	○	3	○	0	x	8	+	6	○	11	○	58	+
6	10	+	19	+	3	○	3	○	0	○	0	x	16	+	3	○	45	+
7	5		5	○	30	+	0	○	5	○	5	○	0	x	9	○	42	+
8	9	+	4	○	14	+	12		18	+	2	○	10		0	x	30	+
9	7	+	7	0	7	○	14	+	21	+	12	+	14	+	19	+	0	x

The probabilities given here are the frame-to-frame transition probabilities as a percentage. Interframe interval is 2 ms. The transition is always from the row PM to the column PM. Pluses indicate a significantly (P<0.5) higher chance of changing into this PM than the *a priori* chance, given by the distribution of centroids. A circle shows that chance of changing into this PM is significantly below the *a priori* chance. Example: the transition probability from hovering (PM9) to an upward flight (PM5) is given in row 9 column 5 and amounts to 21%. It is marked with a plus, so this transition is above chance level. The small x marks omitted transitions. If there is no sign, a gap, this means that the transition is at chance level and not significantly above or below.

Table 3. Transition probabilities as a percentage of the prototypical movements from flights in the large arena

PM	1		2		3		4		5		6		7		8		9	
1	0	x	2	○	26	+	2	○	3	○	5		40	+	9	○	14	○
1	0	x	0	○	21	+	2	○	7	○	14		26	+	18	○	12	○
2	0	○	0	x	4	○	21	+	2		43	+	5	○	9	○	16	
3	3		3	○	0	x	14	+	3	○	0	○	29	+	26	+	23	○
4	0	○	10	+	13	+	0	x	0	○	23	○	0	○	27	+	27	+
5	5	○	5	+	10	○	0	○	0	x	5	+	10	○	5	○	62	+
6	3	+	9	+	0	○	23	○	0	○	0	x	17	+	20	○	29	+
7	9		3	○	22	+	0	○	0	○	19	○	0	x	25	○	22	+
8	7	+	10	○	9	+	10		1	+	20	○	18		0	x	27	+
9	3	+	5	○	14	○	14	+	5	+	19	+	14	+	27	+	0	x

The probabilities given here are the frame-to-frame transition probabilities as a percentage. See Table 2 for details. Pluses indicate a significantly (P<0.5) higher chance of changing into this PM than the *a priori* chance. A circle shows that chance of changing into this PM is significantly below the *a priori* chance. The small x marks omitted transitions. If there is no sign, a gap, this means that the transition is at chance level and not significantly above or below.

with the exception of PM6. Although the individual PMs are to some extent quantitatively different when compared with the corresponding PM in the other arena, they still describe the same type of movement (see Fig. 4).

Irrespective of the environment, the PMs can be classified into rotational PMs and PMs that are characterised by only negligible rotational velocities. This finding is similar to those of previous studies on the blowfly *Calliphora vicina*, which classified behaviour into saccadic and intersaccadic intervals (Schilstra and van Hateren, 1999). Saccades seem to be very distinct prototypical behavioural components. The duration of the saccadic PMs (1, 2, 8) have less variance than the others. Their outer distances calculated in order to determine the quality measures (see Table 1) show that they are quite distinct from other PMs in comparison to the intersaccadic PMs. However, saccades have a large inner variance which led to the detailed analysis (see below). In contrast to saccades, the intersaccadic intervals are much more heterogeneous and could be classified into six PMs, five of which include translational movements.

Characterising sequences of PMs

What are the rules by which complex behaviour is built from PMs? To characterise the order of PMs, we firstly analysed the transition

probabilities between PMs (see Tables 2 and 3), employing a level-one Markov analysis. The transition probabilities from one PM into all other PMs can be used as the stochastic rule that defines how to elongate a behavioural sequence after that PM. We omitted the staying probability for a PM, because this is equivalent to the mean duration of a PM.

We tested whether the transition probabilities were significantly above or below chance level. The null hypothesis was determined by taking into account the a priori probability of the individual PMs. A chance transition into an often-exhibited PM is much more probable than a transition into a rare PM. Fig. 5 shows that for the small arena about 89% (64 out of 72) and for the large arena about 67% (48 out of 72) of the transitions were significantly different from chance level. Some of the 'prohibited' transitions are likely to be a consequence of physical constraints. For example, the transition from a left saccade to a right saccade without an intermediate state is hardly possible for the fly, as it would require an almost instantaneous change of yaw velocity from approximately $+1000\,\mathrm{deg\,s^{-1}}$ to $-1000\,\mathrm{deg\,s^{-1}}$. However, there are also transitions that appear possible for physical reasons but nonetheless do not occur. In the large arena, going from a climb flight to a forward flight is uncommon and *vice versa* (see Fig. 5B, PM5 and PM6). In fact, transitions from climb flight into any forward translational movement seem to be below or at chance level in either arena. In the small arena we can observe that PM9 (hovering) often follows other movements. Five other PMs are significantly more likely to change into the hovering PM than chance level would predict. There is no other PM that is targeted more often significantly above chance level.

Do the transition probabilities of PMs depend on the volume of the flight arena? To answer this question we used Bernoulli statistics and calculated confidence intervals to test whether the transition probabilities from a given PM to another differ significantly in the two flight arenas. We found that 56% of the transitions were not significantly different ($P<0.05$) from each other. When we subtracted the transitions from and to PM6, which is a qualitatively different PM for the two arenas and will elicit different transition probabilities, 60% of the transitions were not significantly different.

When we interpret the transition probabilities as probabilistic rules, we are able to form sequences of PMs, which we call superprototypes. Therefore we omit the staying probabilities and concentrate on the transition probabilities between PMs. An obvious way to form superprototypes is a so-called maximum walk, where one uses the maximal transition probabilities. If we start for example in the first cluster, the superprototype PM1–PM7–PM9 is the maximum walk for the transitions of the small arena. This movement would describe a left saccade (PM1) followed by leftward movement (PM7) and end in hovering (PM9). To test whether the superprototype, created by the maximum walk, is more probable than a chance transition, we used as the null hypothesis the probability given by the product of the independent transitions p. The transition p from a start PM to a target PM is given by the percentage of data assigned to the target PM. The transition probability of the null hypothesis corresponding to the example superprototype is 3.18×10^{-2}. The transition probability based on transitions between the corresponding PMs is 1.67×10^{-1}.

In the aforementioned example, all transitions were significantly above chance level. We used the significant transitions above chance level to create two other examples, shown in Fig. 6. Superprototype A reflects a retreat from a position occupied before (PM6), a saccadic yaw turn to the right (PM2), a short sequence of translational flight (PM4) in the new direction and then hovering (PM9). This combination was found eight times in our data. It is also more likely (2.89×10^{-2}) than the null hypothesis predicts (1.31×10^{-3}).

Another superprototype has already been described qualitatively in hoverflies and blowflies, i.e. zigzagging or wobbling (Collett, 1980; Collett and Land, 1975a; Schilstra and van Hateren, 1999). This superprototype (Fig. 6B) is more likely (7.26×10^{-3}) than the null hypothesis (2.68×10^{-3}) predicts and seems to be a typical movement of flies. It is, for example, consistently performed when flies fly in an elongated tunnel and might be used to analyse the 3D structure of the fly's environment (R.K., unpublished observations).

In conclusion, we derived an objective description of prototypical flight movements and the transition probabilities of their ordering. By applying the clustering procedure we excluded any observer bias.

Fine structure of saccades

All nine PMs show variance. To further elucidate the role of PMs in a behavioural context, we analysed whether this variance is noise overlaying a fixed pattern or systematic variation caused by other flight parameters, as for example body orientation. We concentrated on the rotational PMs, i.e. the so-called saccades. The most prominent change is in yaw angle. Changes in roll or pitch are rather

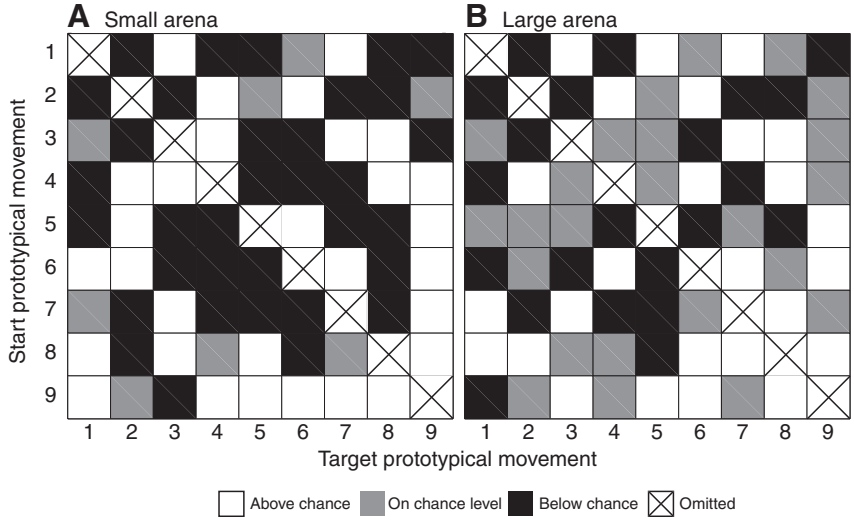

Fig. 5. Significant transitions between body prototypical movements. Transition significance (white, black and grey boxes, see key; $P>0.05$) of the fly changing from one PM to another (A, small arena; B, large arena). Rows mark the starting PM, columns the target PM. Significances were derived by calculating the confidence interval of the transition probability. If the null hypothesis was inside the confidence interval the transition was noted as not significant. The chance level is given by the probability distribution for the occurrence of the target PM. The transition probability values can be found in Tables 2 and 3.

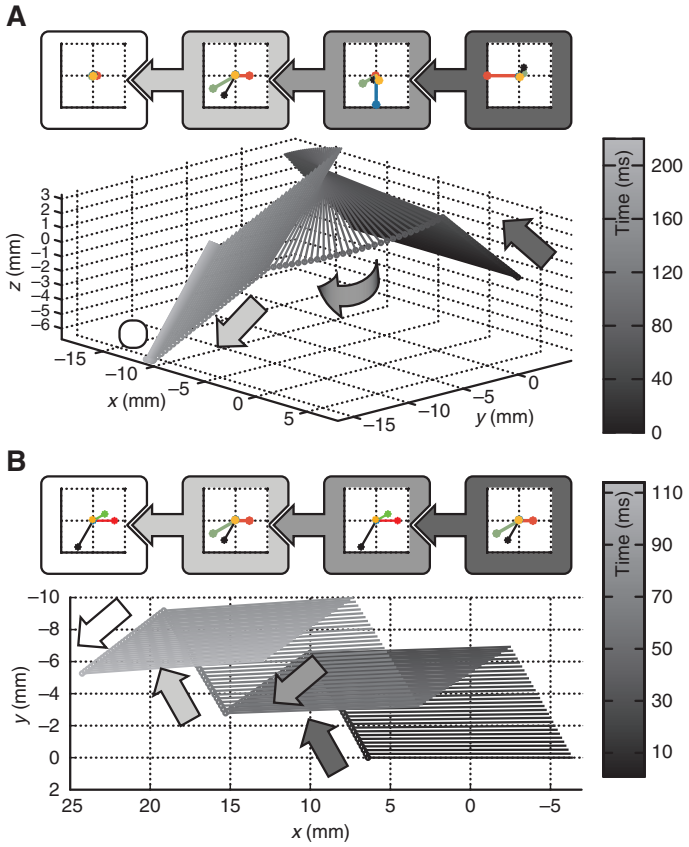

Fig. 6. Superprototypes. (A) Top: an oft-flown sequence of PMs constituting a typical turn of *Eristalis*: backwards movement, right turn, forward–right movement and hovering (from right to left). Bottom: reconstruction of the trajectory flown. The fly is depicted as a line for the body long axis. A circle marks the head position. The grey scale codes for time (see side bar). The arrows mark the phases of movement in the reconstruction. The oval at the end marks the hovering phase. (B) Another superprototype: right–forward, left–forward, right–forward and left–forward. Conventions as in A.

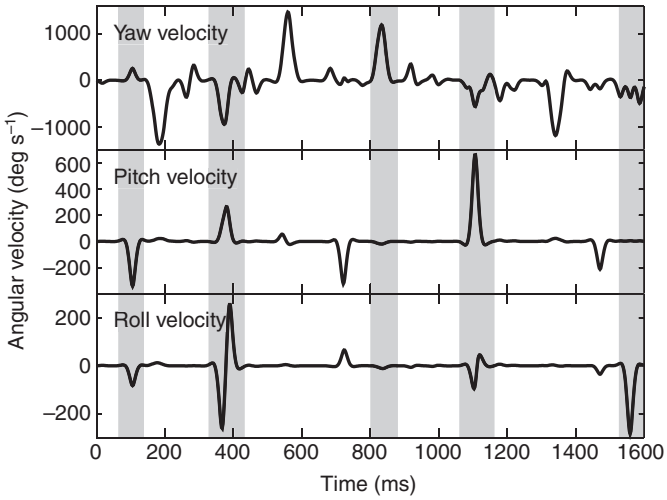

Fig. 7. Coinciding and independently occurring saccades. Head yaw velocity during 1.6 s of a flight trajectory (top row), head pitch velocity (middle row) and head roll velocity (bottom row). Saccade threshold was 400 deg s^{-1} for yaw, 150 deg s^{-1} for pitch and 75 deg s^{-1} for roll velocities. The five saccades highlighted by grey frames are examples of different combinations of saccades.

subtle, but even these delicate changes in head and body orientation are done in a saccadic fashion. Because we wanted to analyse the smaller pitch and roll saccades and compare head and body saccades we had to detect saccades by a thresholding operation with smaller thresholds than the average rotation velocity in PMs (for details of identification see Materials and methods).

We traced head saccades for all three rotational degrees of freedom and found that pitch, yaw and roll saccades do not always coincide. In the short section of a sample trajectory (see Fig. 7) the head performs all combinations of saccades. The mean velocity values of all saccade types are presented in Fig. 8. We combined saccades of opposite direction by taking the absolute values of the rotational velocities and aligning their maximums. Clearly saccades incorporating yaw rotation dominate in terms of number (Fig. 8H). Interestingly, the standard deviation of the mean velocities is quite low for saccades including simultaneous rotations about all three rotational axes (Fig. 8A) and for saccades with only a single rotational axis (Fig. 8E–G). Also the velocities in roll–pitch combination saccades have a relatively small standard deviation (Fig. 8D). The velocities in combination saccades including yaw and only one other rotational axis have large standard deviations and unclear shapes (Fig. 8B,C).

We used the yaw saccades as a trigger to analyse the head and body rotations around the other axes (Fig. 9). Body and head yaw saccades of *Eristalis* have a mean amplitude of 33±22 deg for the body and 35±21 deg for the head. The head yaw has steeper rises and fall-offs than the body yaw (Fig. 9B). These characteristics of *Eristalis* yaw saccades are in accordance with blowfly yaw saccades (van Hateren and Schilstra, 1999). However, the velocity profile of yaw saccades in hoverflies shows a subtle difference to that of blowflies. It seems that *Eristalis* stops turning the head and body simultaneously (Fig. 9D), whereas the head of *Calliphora* has already stopped rotating before the body (van Hateren and Schilstra, 1999). Qualitatively, the other rotational head velocity profiles look like those of *Calliphora*. The change in pitch is smaller than the change in yaw, and the roll movements are even more subtle than pitch rotations (see Fig. 9C). During rotations as well as during forward–sideways flight (compare PM3 and PM4) we found dramatic changes in body roll ranging up to approximately 80 deg (estimated by eye inspection of the high-speed image frames). The maximum roll angle of the head was, however, only 12 deg during a saccade. Thus, the head roll angle is stabilised quite well against body roll, as is also characteristic of *Calliphora* (van Hateren and Schilstra, 1999; Hengstenberg et al., 1986; Schilstra and van Hateren, 1999).

How are direction and amplitude of yaw saccades controlled? Collett and Land suggested the angle between body long axis and course direction (ψ angle) had an impact on the direction of yaw saccades in the hoverfly *Syritta pipiens* (Collett and Land, 1975a) as did Wagner for houseflies (Wagner, 1986). We analysed this ψ angle in our trajectories and found it to vary much more in *Eristalis* than in houseflies (Wagner, 1986). The ψ angles range from 0 to 180 deg in *Eristalis*. A similar relationship is found in *Eristalis* to that in *Syritta* between ψ and saccade direction: saccades are often directed to reduce ψ (Fig. 10). Moreover the ψ amplitude resulting from the yaw saccade correlates with the starting ψ angle. Larger starting ψ angles correlate with larger ψ amplitudes (see Fig. 10). This finding shows that the ψ angle is one source of the variance in yaw saccades as represented by PM1 and PM2.

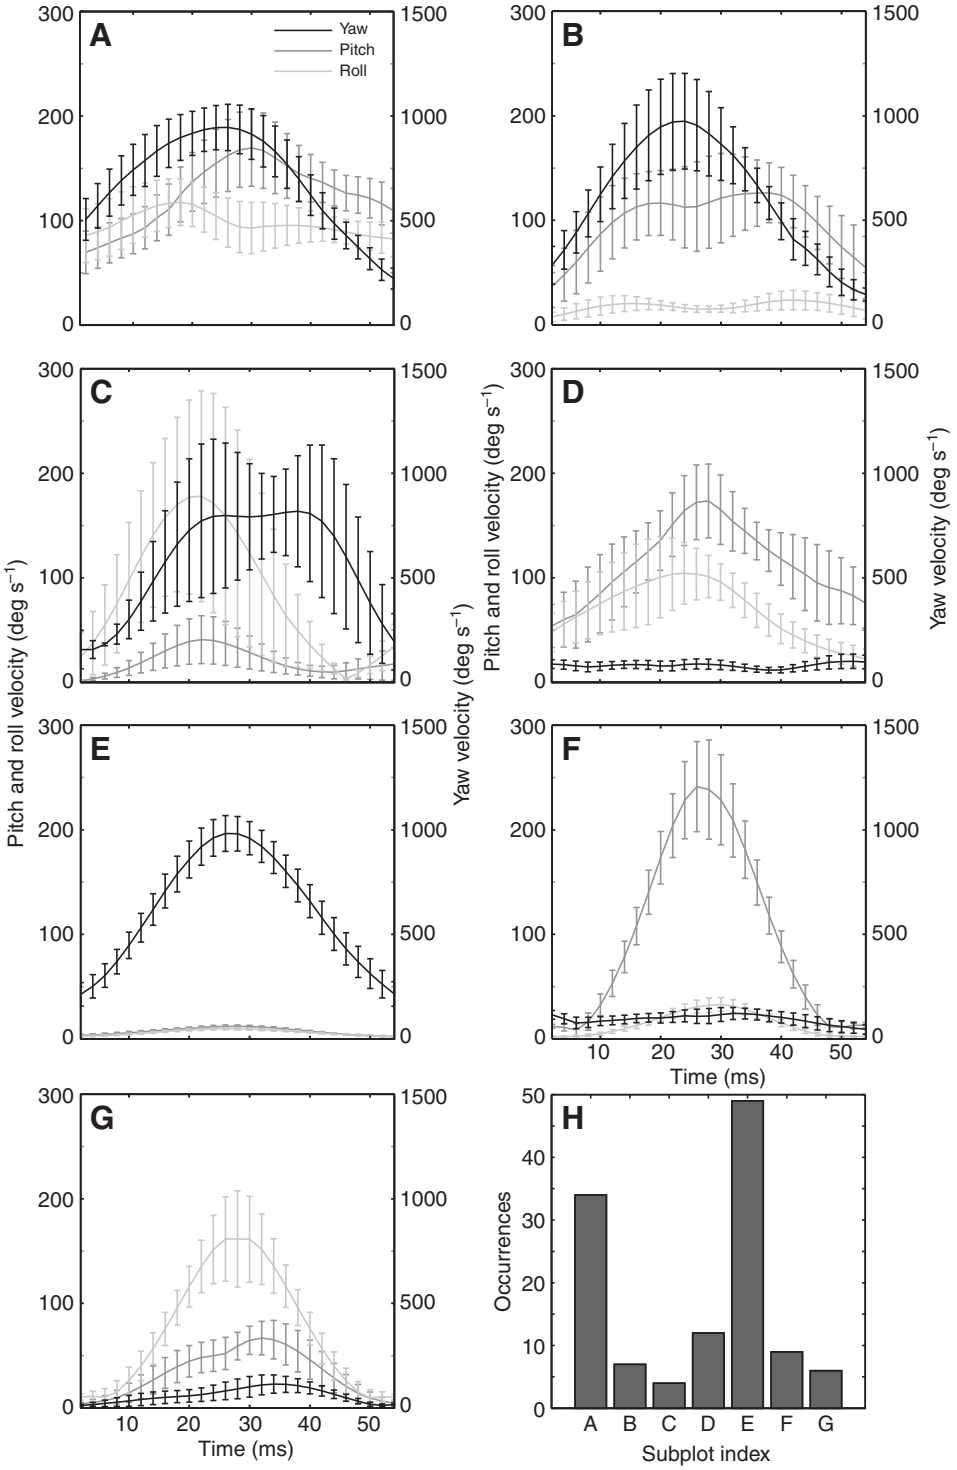

Fig. 8. Different head saccade types. (A–G) Mean angular velocities (±s.e.m.) during saccades plotted against time. (A) Yaw saccades coinciding with pitch and roll saccades (*N*=34). (B) Yaw saccades coinciding with pitch saccades (*N*=7). (C) Yaw saccades coinciding with roll saccades (*N*=4). (D) Coinciding pitch and roll saccades (*N*=12). (E) Pure yaw saccades (*N*=49). (F) Pure pitch saccades (*N*=9). (G) Roll saccades (*N*=6). (H) Number of head saccade occurrences. *x*-axis denotes type by referring to the plots A–G.

As is characteristic of the head yaw angle, the head pitch jitters less and shows longer periods of stable orientation than the body (Fig. 11A). The changes in head pitch are steeper and the pitch angle of the head is at all times smaller than that of the body. Since body pitch velocities have a much wider distribution, it is suggested that the head is stabilised against most pitch movements of the body. In our head data set we counted 113 head pitch saccades *versus* 532 body pitch saccades.

To explain the difference in the number of body and head pitch saccades and the rather weak correlation between body and head pitch (covariance normalised by the product of the autocorrelation: 0.46±0.07 mean ± s.d.; data not shown), we analysed the relationship between ground velocity and body pitch and head pitch (Fig. 11B). We calculated the covariance for both, normalised by the autocorrelation between ground speed and head pitch and body pitch, respectively. The normalisation was done because pitch and ground velocity have different dimensions and numerical magnitudes. Every flight body trajectory of 4 s length, for which a head trajectory was also available (*N*=5) was analysed in this way. The mean covariance value and standard deviation were derived from individual covariance analyses. The correlation between head and body pitch is still stronger than the correlation between head pitch

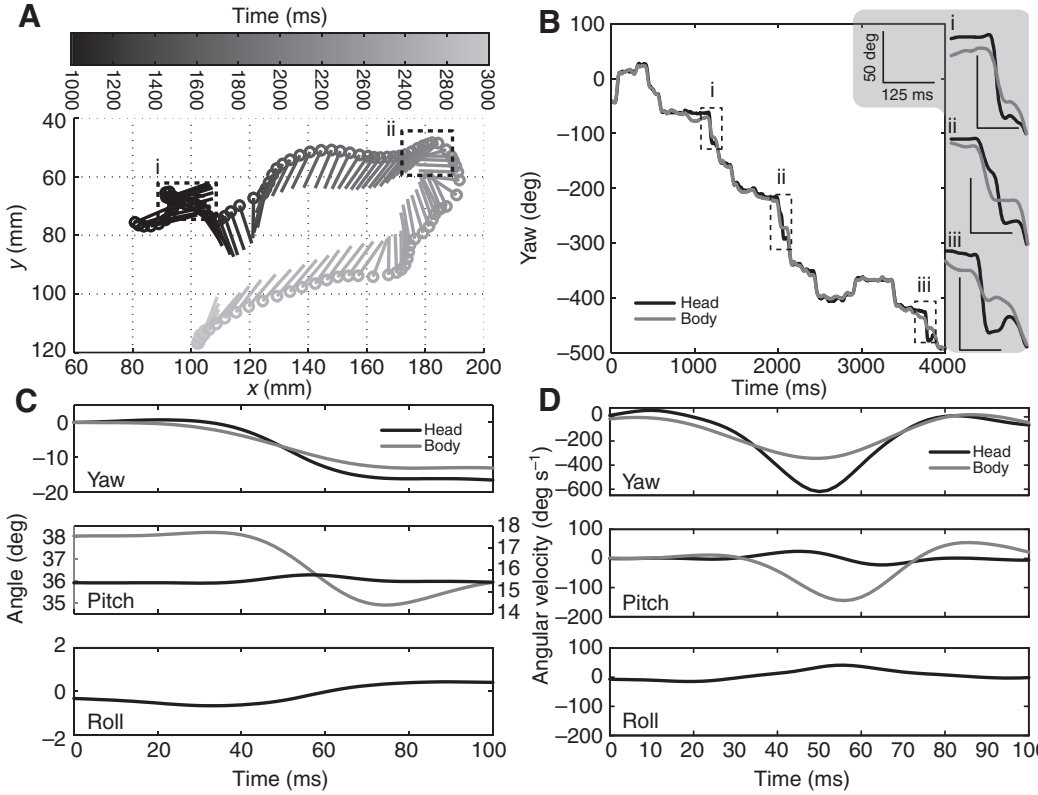

Fig. 9. Yaw saccades. (A) A 3 s example flight trajectory of *Eristalis* seen from above. The circles denote head positions, lines mark the body long axis (time colour coded, see colour bar). Dashed boxes (i,ii) mark yaw saccades. (B) Head and body yaw angles during the example flight in A plotted against time. The angle is plotted for the complete trajectory (4 s). Dashed boxes (i–iii) mark saccades, shown in detail on the right (scale bars in top right of figure). The third saccade is not shown in the trajectory in A. (C) Mean head and body absolute angles during typical saccades (N=42). Saccades of both directions were used. Starting yaw value is set to zero, because the starting yaw orientation is arbitrary for the saccade. Note second y-axis for pitch angles; the axis on the left is the body pitch axis; the axis on the right is the head pitch axis. (D) Mean angular velocities of head and body during saccades.

and ground velocity, which is another indication that the head is stabilised against these body movements. In *Eristalis* the body pitch is negatively correlated to ground velocity. This was also found to be the case in houseflies (Wagner, 1986) and fruit flies (David, 1978). The body pitch angle of *Eristalis* is steepest during backward flight (Fig. 11C). Hence the pitch angle and pitch saccades are also influenced by other flight parameters, for example the ground speed.

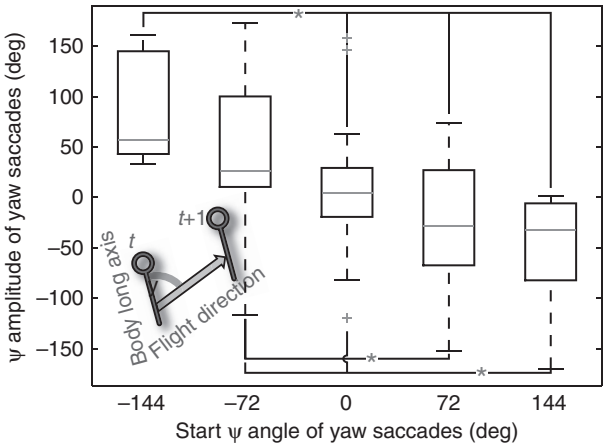

Fig. 10. Influence of the ψ angle on yaw saccades. The ψ amplitudes of yaw saccades are binned according to the ψ angle at the start of the saccade (see x-axis labels, bins are 72 deg wide). The ψ amplitudes of the yaw saccades are plotted as box–whisker plots. The grey midline in each box marks the median of the respective data. The upper and lower parts of the box show the upper and lower quartile. The whiskers denote 1.5 times the interquartile range, which is the distance from top to bottom of each box. Points lying outside 1.5 times the interquartile range are plotted as grey crosses. The populations were tested with an ANOVA; significantly different bins are marked by an asterisk. The pictogram at the bottom left corner illustrates the ψ angle (the angle between the body long axis and the flight direction, where t is time).

A roll rotation of the head is observed only rarely as *Eristalis* stabilises the head almost perfectly against body roll movements. During intersaccadic intervals the mean absolute roll velocity is 1.2±2.2 deg s^{-1}. Nonetheless, head roll saccades are performed as can be seen in Fig. 7. As a first approach to find the causes of these saccades, we classified head roll saccades as either increasing (N=24) or decreasing the roll angle (N=35). We assumed, based on the roll angle of the intersaccadic interval, that the fly stabilises its head horizontally. We analysed the roll angle at the onset and end of saccades and found that the start angles of saccades which decrease the roll angle are significantly larger than the median intersaccadic roll angle (see Fig. 12). These saccades turn the head to approximately its normal horizontal orientation and are therefore called correction saccades. The end roll angles of correction saccades are not significantly different from the median roll angle of all intersaccadic intervals. Interestingly, non-correction saccades often turn the head to an orientation above the median intersaccadic roll angle, but start at angles that are similar to the intersaccadic median.

We analysed the body roll of the fly during the 24 non-correction saccades by eye. Eleven saccades (46%) were performed in the opposite direction to concurrent body roll movements. Six saccades (25%) were directed in the same direction as the body roll. Two roll saccades (8%) were found during landing and during a yaw saccade without any body roll. In the final 21% (N=5) we could not find any indication of body roll.

DISCUSSION

We categorised complex flight behaviour of the hoverfly, *E. tenax*, under spatially constrained conditions into prototypical movements (PMs), and analysed further details about a particular kind of PM, the saccades, which form a distinguishing feature of the behavioural

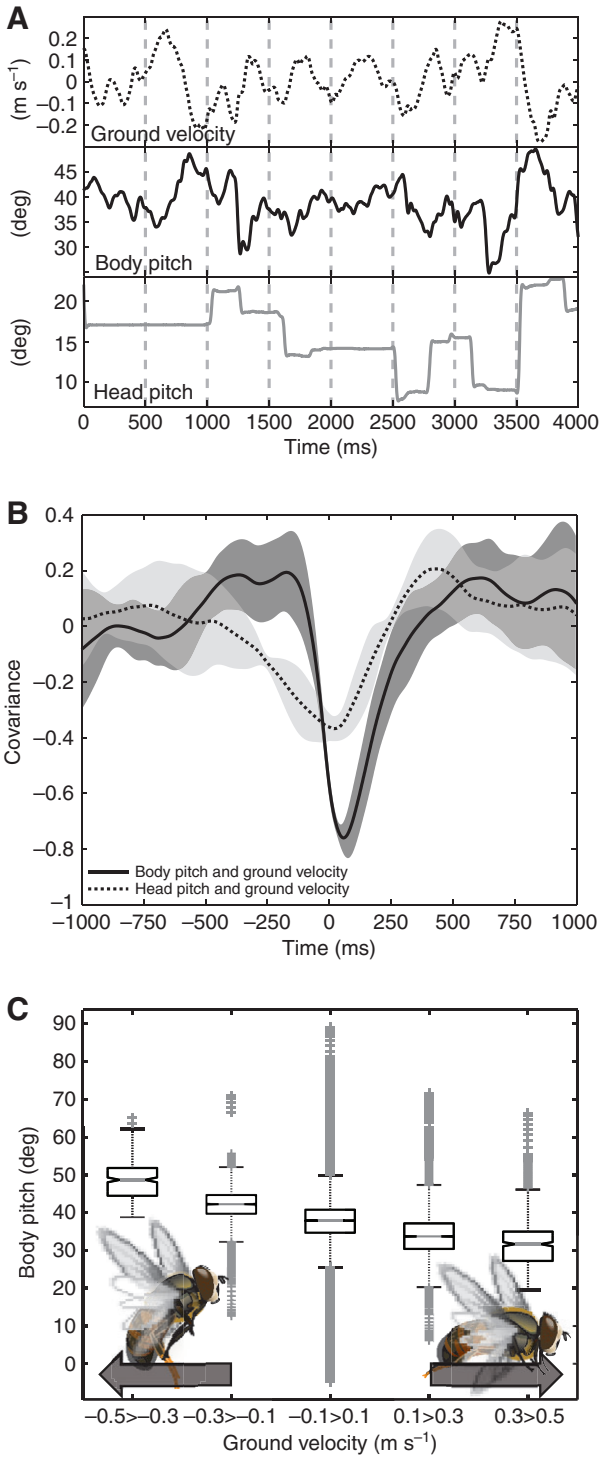

Fig. 11. Pitch dependencies. (A) Ground velocity, body pitch and head pitch of a hoverfly during an example flight plotted against time. (B) Mean (±s.d.) covariance between ground velocity and head or body pitch. The covariance was normalised by the respective products of the autocorrelograms. The covariance was calculated for every flight. We calculated the mean covariance and standard deviation of the mean from all flights of at least 4 s length (N=5). (C) The ground velocity was binned according to the x-axis. Box–whisker plots show the body pitch angle. The grey horizontal line displays the median. The upper and lower boxes represent the upper and lower quartile. The distance between top and bottom of one box is its interquartile range. The whiskers mark 1.5 times that range. Data points outside 1.5 times the interquartile range are plotted as grey crosses. Large arrows indicate the direction of flight.

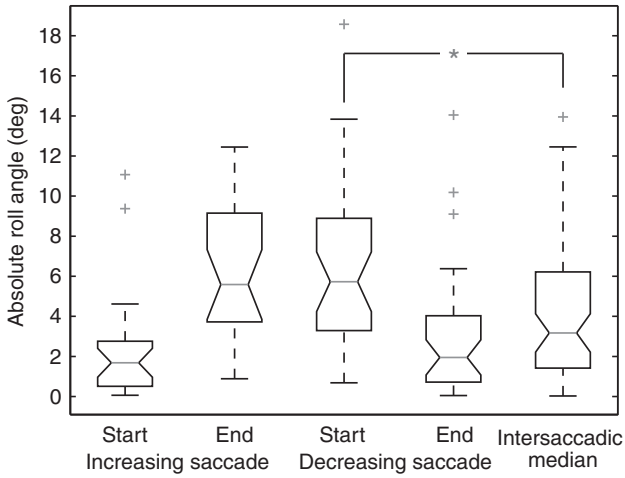

Fig. 12. Roll saccades. From left to right, the first two boxplots show the absolute start and end angle of saccades increasing the roll angle (N=24). The third and fourth boxplots show the start and end angles of roll-decreasing saccades, termed correction saccades (N=35). The fifth boxplot represent the median roll angle in the intersaccadic intervals. All boxplots are organised as described in Fig. 10. Non-overlaying box notches mark distributions with significantly different median (P=0.05). For better visualisation, we included an asterisk to mark that the start angle of saccades decreasing the roll angle are significantly higher than the intersaccadic median roll angle. All other groups do not differ significantly from the intersaccadic median.

repertoire of many insects. By applying a customised k-means clustering analysis to the normalised behavioural velocity data we derived nine PMs from continuous flight trajectories, which constitutes a tremendous complexity reduction of behavioural description. Furthermore, we analysed the ordering of PMs and established a Markov model, giving insight into the way behaviour is organised. Although first-order transitions of a Markov model are unlikely to reveal the complete organisation of flight behaviour, it is a good tool with which to find reoccurring successions of flight manoeuvres, as we explain in detail below. This model together with the set of prototypes constitutes a compact ethogram for *Eristalis* for the confined flight environments and conditions of our experiments.

The nine identified PMs are the result of applying a general purpose data-mining technique with the k-means algorithm. Thereby, we had to make several assumptions concerning feature selection, normalisation, similarity measurement, clustering procedure and evaluation. Each of them generally influences the results. However, some of the assumptions we made are constrained by prior knowledge; others are used for the reason of computational simplicity. At the feature selection step we choose the features that are in common use for describing movements, the fly-centred velocities. However, we additionally tested an alternative world-based velocity feature set that rendered qualitatively the same results. For normalisation we applied the standard z-score approach in order to ensure that each velocity dimension contributes equally to the clustering. Using the k-means with squared Euclidean distances for classifying the velocity data also constitutes a widely used approach because of its robustness against noise and computational simplicity. Within the k-means clustering approach we selected the most appropriate classifier and evaluated it by comparing the sets of classes resulting from determining different numbers of clusters. For this evaluation we applied general criteria for an appropriate clustering.

After determining the set of nine classes for our data, we were able to classify each point of a flight trajectory and to determine the temporal sequence of PMs. Analysing this sequence with respect to the duration of the occurring PMs allowed us to assess whether the classification leads to a meaningful segmentation of the trajectory. The mean durations of PMs take values in the range of several tens of milliseconds. Nonetheless, we also found much shorter durations. These we treated as classification errors, because owing to physical constraints genuine movements of hoverflies can be expected to be in the range of more than 10 ms. Classification errors occur if a behavioural sequence does not fit well into one class but, instead, is located at the border between several classes. Also, the exact transitions between PMs (where the saccade ends and where the sideways motion begins) might be smooth and result in false classifications for the transition region. This form of classification noise is included in our data and thus also occurs as noise for the calculation of the transition probabilities between PMs.

Furthermore, we tested how PMs may depend on a 21-fold change in the volume of the flight arena. Only one PM changed between these different conditions. The transition probabilities between PMs in the different flight arenas stayed similar as well. Nonetheless, we would expect changes in the set of PMs if the environment allowed for entirely new behavioural aspects. Barriers or defined landing sites, for example, are likely to elicit different behavioural components, and therefore to lead to the occurrence of other PMs with other transition probabilities. Also, much larger flight arenas than were used in this study will allow the animal to attain higher speeds.

The ethogram obtained in this way for *Eristalis* firstly shows a separation into rotational and translational movements similar to those described for other flying species (Bender and Dickinson, 2006; Boeddeker and Hemmi, 2010; Collett and Land, 1975a; Eckmeier et al., 2008; van Hateren and Schilstra, 1999; Mronz and Lehmann, 2008; Schilstra and van Hateren, 1999; Wagner, 1986). As a consequence of this behavioural segregation into translational and rotational movements, the translational optic flow generated on the eyes between saccades is not contaminated by much rotational flow. This feature might be of computational significance, because only the translational optic flow component contains information about the three-dimensional layout of the environment (Koenderink, 1986). In blowflies it was shown that motion-sensitive tangential cells represent spatial information during the intersaccadic intervals (Boeddeker et al., 2005; Egelhaaf, 2006; Karmeier et al., 2006; Kern et al., 2005; Kern et al., 2006; Lindemann et al., 2005). Whereas saccades are reflected by just a pair of distinct PMs, the intersaccadic interval, interestingly, is subclassified into six PMs, five of which include significant translational movements.

During saccades, the gaze strategy and head–body coordination of *Eristalis* are similar to those of *Calliphora* (van Hateren and Schilstra, 1999; Schilstra and van Hateren, 1999). Both flies minimise the duration of rotational optic flow on the retina. In both cases, head saccades have a larger velocity amplitude and shorter duration than body saccades, which further reduces the duration of rotational optic flow. However, in contrast to *Calliphora* where the head stops rotating before the body, the head of *Eristalis* terminates its rotation at the same time as the body does. We observed *Eristalis* to exhibit all possible combinations of rotations in a saccadic way.

The pitch and yaw saccades of the head might be used to fixate new targets: roll saccades, instead, do not seem to be necessary to change the gaze towards a new target. Motion vision cells like the tangential cells of flies respond predominantly to either horizontal or vertical motion. These cells are often regarded as matched filters for the detection of self-motion (Karmeier et al., 2006; Krapp et al., 1998; Krapp et al., 2000). Hence, stabilisation of head roll orientation would simplify the later processing of optic flow information by the visual pathway. Without this stabilisation, the tangential cells, for example, would respond to a yaw rotation with a horizontally aligned head in the same way as to a pitch rotation with a vertically aligned head. Our findings show that the head is stabilised with minimal drift during intersaccadic intervals. There are two subclasses of roll saccades. One class corrects the head position back to a horizontal orientation, which is in accordance with the need to stabilise the head against roll movement of the body. The other class corresponds to head turns away from the horizontal orientation. They are in most cases associated with body roll movements, either in the opposite or the same direction as the body roll. These head roll saccades might be either body roll residuals or overcompensation by the head.

Although the set of PMs is rather stable against changes in the size of the flight arena, every PM is highly variable. We analysed whether this variability is pure noise in a fixed motion pattern or a result of systematic adjustments of PMs to flight parameters. We found that the orientation of the body relative to the flight trajectory (ψ angle) influences the direction and amplitude of a saccade (PM1, PM2) as has already been shown for other flies (Collett and Land, 1975a; Wagner, 1986). It seems that although hoverflies can fly in any direction, they often align their body with the flight direction. This is an indication that PM1 and PM2 are adjusted to the ψ angle. We also found that the body pitch is negatively correlated to ground velocity, as was shown for other flies (David, 1978; Wagner, 1986). This behaviour might reduce friction during faster flight. The body pitch angle is even steeper while flying backwards than during any other observed behaviour. The head is stabilised fairly well against body pitch movements, but there are weak correlations between body and head pitch as well as between head pitch and ground velocity. These might be residual effects of the body pitch. Both yaw and pitch saccades correlate systematically with other flight parameters, which show that PMs are adjusted to the actual flight situation.

What are the rules behind the arrangement of PMs as the building blocks of complex behaviour? A similar analysis to that done here was undertaken in the face-grooming behaviour of mice (Fentress and Stilwell, 1973). PMs were defined and the transition probabilities between such movements were characterised. We adapted this approach and determined the transition probabilities on a frame-by-frame basis. A structure consisting of building blocks and rules is regarded as syntax in formal languages. To derive the rules or in this context the grammar from the succession of our building blocks, we analysed the transition probabilities of PMs employing a level-one Markov analysis. Markov chains are regarded as an extension of regular grammars (Fu, 1974; Gonzalez and Thomason, 1978). A regular grammar is the set of rules forming a formal language (type 3 grammar) in the Chomsky hierarchy (Chomsky and Schützenberger, 1963).

We employed the syntax of PMs to derive more complex flight behaviour, so-called superprototypes. For the first kind of superprototype we always used the maximal transition probability and derived a short succession of a turn, sideways flight and hovering. For the examples shown in Fig. 6, we used the significant transitions shown in Fig. 5. The most probable transition and therefore the most probable superprototype is not always the most interesting one. If one calculated the word occurrence probability in the English language, based on standard newspaper texts, words like ‘and’ would be far more probable than ‘xylophone’.

For *Eristalis* behaviour the ‘xylophones’, such as the zigzagging superprototype, might be particularly useful in actively probing the spatial structure of the environment. The two PMs involved in

zigzagging (PM3 and PM4) are characterised by pronounced sideways velocities, similar to the peering movements of locusts (Collett and Paterson, 1991; Sobel, 1990; Wallace, 1959) and mantids (Kral and Poteser, 1997). These animals use the resulting motion parallax information to judge distances. In robots this gaze strategy was successfully used to navigate around obstacles (Sobey, 1994). *Eristalis* may employ zigzagging in the same way locusts use peering and thereby acquire more information about the 3D layout of its surroundings. The other superprototype, consisting of a backwards flight, turning, drifting and hovering, might be typical wall-avoidance behaviour and would accord with our hypothesis that the constraints set by the small arena led to more saccades than those set by the large one.

Behavioural analysis based on PMs and their syntax has great potential for interspecies comparisons of behaviour. *Eristalis* has been claimed to mimic foraging honeybees, to be less attractive to potential predators (Golding et al., 2001; Golding and Edmunds, 2000). Foraging bees navigate in close proximity to flowers and only fly slowly for short durations. Similarly, *Eristalis* in a confined space only flew for short periods at low velocities. Thus it might be possible that at low flight speeds, as exhibited in our arenas or during the last phase of homing flights (Collett and Land, 1975b), *Eristalis* mimics the honeybee to avoid being attractive prey. It might be possible to address this hypothesis by comparing the ethograms of bees and hoverflies based on PMs.

LIST OF SYMBOLS AND ABBREVIATIONS

k	number of clusters
PM	prototypical movement
t	time
ϕ	pitch angle
θ	yaw angle
ψ	angle between the body long axis and the flight direction

ACKNOWLEDGEMENTS

We would like to thank the Deutsche Forschungsgemeinschaft (DFG) and the Studienstiftung des deutschen Volkes for supporting this study. We thank Grid Schwerdtfeger for her extensive help during head tracking, Nicole Carey for linguistic improvement of the manuscript and Christian Spalthoff for improving the figures.

REFERENCES

Barnett, P. D., Nordstrom, K. and O'Carroll, D. C. (2007). Retinotopic organization of small-field-target-detecting neurons in the insect visual system. *Curr. Biol.* **17**, 569-578.

Bender, J. A. and Dickinson, M. H. (2006). Visual stimulation of saccades in magnetically tethered *Drosophila*. *J. Exp. Biol.* **209**, 3170-3182.

Boeddeker, N. and Hemmi, J. M. (2010). Visual gaze control during peering flight manoeuvres in honeybees. *Proc. Biol. Sci.* **277**, 1209-1217.

Boeddeker, N., Lindemann, J. P., Egelhaaf, M. and Zeil, J. (2005). Responses of blowfly motion-sensitive neurons to reconstructed optic flow along outdoor flight paths. *J. Comp. Physiol. A* **191**, 1143-1155.

Boeddeker, N., Dittmar, L., Stürzl, W. and Egelhaaf, M. (2010). The fine structure of honeybee head and body yaw movements in a homing task. *Proc. Biol. Sci.* **277**, 1899-1906.

Bouguet, J. Y. (1998). Camera calibration from points and lines in dual-space geometry. Technical report, California Institute of Technology.

Braun, E., Geurten, B. and Egelhaaf, M. (2010). Identifying prototypical components in behaviour using clustering algorithms. *PLoS ONE* **5**, e9361.

Chomsky, N. and Schützenberger, M. P. (1963). The algebraic theory of context free languages. In *Computer Programming and Formal Systems* (ed. P. Braffort and D. Hirschberg), pp. 118-161. Amsterdam: North Holland Publishing Company.

Collett, T. S. (1980). Some operating rules for the optomotor system of a hoverfly during voluntary flight. *J. Comp. Physiol. A* **138**, 271-282.

Collett, T. S. and Land, M. F. (1975a). Visual control of flight behaviour in the hoverfly *Syritta pipiens* L. *J. Comp. Physiol. A* **99**, 1-66.

Collett, T. S. and Land, M. F. (1975b). Visual spatial memory in a hoverfly. *J. Comp. Physiol. A* **100**, 59-84.

Collett, T. S. and Land, M. F. (1978). How hoverflies compute interception courses. *J. Comp. Physiol. A* **125**, 191-204.

Collett, T. S. and Paterson, C. J. (1991). Relative motion parallax and target localisation in the locust, *Schistocerca gregaria*. *J. Comp. Physiol. A* **169**, 615-621.

David, C. T. (1978). Relationship between body angle and flight speed in free-flying *Drosophila*. *Physiol. Entomol.* **3**, 191-195.

Eckmeier, D., Geurten, B. R. H., Kress, D., Mertes, M., Kern, R., Egelhaaf, M. and Bischof, H. J. (2008). Gaze strategy in the free flying zebra finch (*Taeniopygia guttata*). *PLoS ONE* **12**, e3956.

Egelhaaf, M. (2006). The neural computation of visual motion information. In *Invertebrate vision* (ed. E. Warrant and D. E. Nielsson), pp. 399-461. Cambridge: Cambridge University Press.

Fentress, J. C. and Stilwell, F. P. (1973). Grammar of a movement sequence in inbred mice. *Nature* **244**, 52-53.

Flash, T. and Hochner, B. (2005). Motor primitives in vertebrates and invertebrates. *Curr. Opin. Neurobiol.* **15**, 660-666.

Fu, K. S. (1974). *Syntactic Methods in Pattern Recognition.* New York, USA: Academic Press Inc.

Golding, Y. C. and Edmunds, M. (2000). Behavioural mimicry of honeybees (*Apis mellifera*) by droneflies (Diptera: Syrphidae: Eristalis spp.). *Proc. R. Soc. Lond. B. Biol. Sci.* **267**, 903-909.

Golding, Y. C., Ennos, A. R. and Edmunds, M. (2001). Similarity in flight behaviour between the honeybee *Apis mellifera* (Hymenoptera: apidae) and its presumed mimic, the dronefly *Eristalis tenax* (Diptera: syrphidae). *J. Exp. Biol.* **204**, 139-145.

Gonzalez, R. C. and Thomason, M. G. (1978). *Syntactic Pattern Recognition.* Massachusetts, USA: Addison-Wesley Publishing Company.

Graziano, M. S. A., Taylor, C. S. R. and Moore, T. (2002). Complex movements evoked by microstimulation of precentral cortex. *Neuron* **34**, 841-851.

Graziano, M. S. A., Aflalo, T. N. S. and Cooke, D. F. (2005). Arm movements evoked by electrical stimulation in the motor cortex of monkeys. *J. Neurophysiol.* **94**, 4209-4223.

Hengstenberg, R., Sandeman, D. C. and Hengstenberg, B. (1986). Compensatory head roll in the blowfly *Calliphora* during flight. *Proc. R. Soc. Lond. B. Biol. Sci.* **227**, 455-482.

Jenkins, O. C. and Mataric, M. J. (2003). Automated derivation of behavior vocabularies for autonomous humanoid motion. *Proceedings of the Second International Joint Conference on Autonmous Agents and Multiagent Systems,* pp. 225-232. New York, NY, USA: Association for Computing Machinery.

Karmeier, K., van Hateren, J. H., Kern, R. and Egelhaaf, M. (2006). Encoding of naturalistic optic flow by a population of blowfly motion-sensitive neurons. *J. Neurophysiol* **96**, 1602-1614.

Kern, R., van Hateren, J. H., Michaelis, C., Lindemann, J. P. and Egelhaaf, M. (2005). Function of a fly motion-sensitive neuron matches eye movements during free flight. *PLoS Biol.* **3**, 1130-1138.

Kern, R., van Hateren, J. H. and Egelhaaf, M. (2006). Representation of behaviourally relevant information by blowfly motion-sensitive visual interneurons requires precise compensatory head movements. *J. Exp. Biol.* **209**, 1251-1260.

Koenderink, J. J. (1986). Optic flow. *Vis. Res.* **26**, 161-179.

Kral, K. and Poteser, M. (1997). Motion parallax as a source of distance information in locusts and mantids. *J. Insect Behav.* **10**, 145-163.

Krapp, H. G., Hengstenberg, B. and Hengstenberg, R. (1998). Dendritic structure and receptive-field organization of optic flow processing interneurons in the fly. *J. Neurophysiol.* **79**, 1902-1917.

Krapp, H. G., Dahmen, H. J. and Franz, M. O. (2000). Extracting ego-motion from optic flow: limits of accuracy and neuronal filters. In *Motion Vision: Computational, Neural and Ecological Constraints* (ed. J. M. Zanker and J. Zeil), pp. 143-168. Berlin: Springer.

Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* **2**, 83-97.

Kuhn, H. W. (1956). Variants of the Hungarian method for assignment problems. *Naval Research Logistics Quarterly* **3**, 253-258.

Liang, R.-H. and Ouhyoung, M. (1998). A real-time continuous gesture recognition system for sign language. *Third IEEE International Conference on Automatic Face and Gesture Recognition*, pp. 558-567.

Lindemann, J. P., Kern, R., van Hateren, J. H., Ritter, H. and Egelhaaf, M. (2005). On the computations analyzing natural optic flow: quantitative model analysis of the blowfly motion vision pathway. *J. Neurosci.* **25**, 6435-6448.

Lindemann, J. P., Weiss, H., Möller, R. and Egelhaaf, M. (2008). Saccadic flight strategy facilitates collision avoidance: closed-loop performance of a cyberfly. *Biol. Cybern.* **98**, 213-227.

MacQueen, J. B. (1967). Some methods for classification and analysis of multivariate observations. *Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability* **1**, 281-297.

Mronz, M. and Lehmann, F. O. (2008). The free-flight response of *Drosophila* to motion of the visual environment. *J. Exp. Biol.* **211**, 2026-2045.

Mussa-Ivaldi, F. A. and Bizzi, E. (2000). Motor learning through the combination of primitives. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **355**, 1755-1769.

Nordström, K., Barnett, P. D., Moyer de Miguel, I. M., Brinkworth, R. S. and O'Carroll, D. C. (2008). Sexual dimorphism in the hoverfly motion vision pathway. *Curr. Biol.* **18**, 661-667.

O'Carroll, D. C., Bidwell, N. J., Laughlin, S. B. and Warrant, E. J. (1996). Insect motion detectors matched to visual ecology. *Nature* **382**, 63-66.

O'Carroll, D. C., Laughlin, S. B., Bidwell, N. J. and Harris, R. A. (1997). Spatio-temporal properties of motion detectors matched to low image velocities in hovering insects. *Vis. Res.* **37**, 3427-3439.

Schilstra, C. and van Hateren, J. H. (1999). Blowfly flight and optic flow I. Thorax kinematics and flight dynamics. *J. Exp. Biol.* **202**, 1481-1490.

Sobel, E. C. (1990). The locust's use of motion parallax to measure distance. *J. Comp. Physiol. A* **167**, 579-588.

Sobey, P. (1994). Active navigation with a monocular robot. *Biol. Cybern.* **71**, 433-440.

Stepniewska, I., Fang, P. C. and Kaas, J. H. (2005). Microstimulation reveals specialized subregions for different complex movements in posterior parietal cortex of prosimian galagos. *Proc. Natl. Acad. Sci.* **102**, 4878-4883.

Tammero, L. F. and Dickinson, M. H. (2002). The influence of visual landscape on the free flight behavior of the fruit fly *Drosophila melanogaster. J. Exp. Biol.* **205**, 327-343.
Thurau, C. and Hlavac, V. (2007). n-grams of action primitives for recognizing human behavior. *Proceedings of the 12th International Conference on Computer Analysis of Images and Patterns (CAIP'07)* **4673**, 93-100.
van der Schaaf, A. and van Hateren, J. H. (1996). Modelling the power spectra of natural images: statistics and information. *Vis. Res.* **17**, 2759-2770.
van Hateren, J. H. and Schilstra, C. (1999). Blowfly flight and optic flow. II. Head movements during flight. *J. Exp. Biol.* **202**, 1491-1500.
Wagner, H. (1986). Flight performance and visual control of flight of the free-flying housefly (*Musca domestica* L.) I. Organization of the flight motor. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **312**, 527-551.
Wallace, G. K. (1959). Visual scanning in the desert locust *Schistocerca gregaria* Forskål. *J. Exp. Biol.* **36**, 512-525.