
Automatic detection and decoding of honey bee waggle dances

Fernando Wario, Benjamin Wild, Raúl Rojas, Tim Landgraf*

FB Mathematik und Informatik, Freie Universität Berlin, Berlin, Germany

* tim.landgraf@fu-berlin.de

Abstract

The waggle dance is one of the most popular examples of animal communication. Forager bees direct their nestmates to profitable resources via a complex motor display. Essentially, the dance encodes the polar coordinates to the resource in the field. Unemployed foragers follow the dancer's movements and then search for the advertised spots in the field. Throughout the last decades, biologists have employed different techniques to measure key characteristics of the waggle dance and decode the information that it conveys. Early techniques involved the use of protractors and stopwatches to measure the dance orientation and duration directly from the observation hive. Recent approaches employ digital video recordings from which dance characteristics are extracted using digital protractors on screen. However, manual approaches are very time-consuming. Most studies, therefore, regard only small numbers of animals in short periods of time. We have developed a system capable of automatically detecting, decoding and mapping communication dances in real-time. In this paper, we describe our recording setup, the image processing steps performed for dance detection and decoding and an algorithm to map dances to the field. The proposed system performs with an accuracy of 90.07%. The decoded waggle orientation has an average error of -2.92° ($SD = 7.37^\circ$), well within the range of human error. To evaluate and exemplify the system's performance, a group of bees was trained to an artificial feeder, and all dances in the colony were automatically detected, decoded and mapped. The system presented here is the first of this kind made publicly available, including source code and hardware specifications. We hope this will foster quantitative analyses of the honey bee waggle dance.

Introduction

The honey bee waggle dance is one of the most popular communication systems in the animal world. Forager bees move in a stereotypic pattern on the honeycomb to share the location of valuable resources with their nestmates [18, 35, 44]. Dances consist of waggle and return phases. During the waggle phase, the dancer vibrates her body from side to side while moving forward in a rather straight line. To each waggle phase follows a return phase, during which the dancer circles back to the starting point of the waggle phase. Right and left return phases are alternated, such that the dancer describes a path resembling the figure eight [27, 40, 44, 46]. The average orientation of successive waggle runs with respect to the vertical in the hive approximates the angle between the advertised resource and the solar azimuth as seen from the hive Fig 1. The duration of

the waggle run correlates with the distance between hive and resource [9, 12, 13, 44]. It has also been found a correlation between the dance's timing and the resource's profitability: highly profitable sites are signaled with shorter return runs, yielding a higher waggle production rate [37].

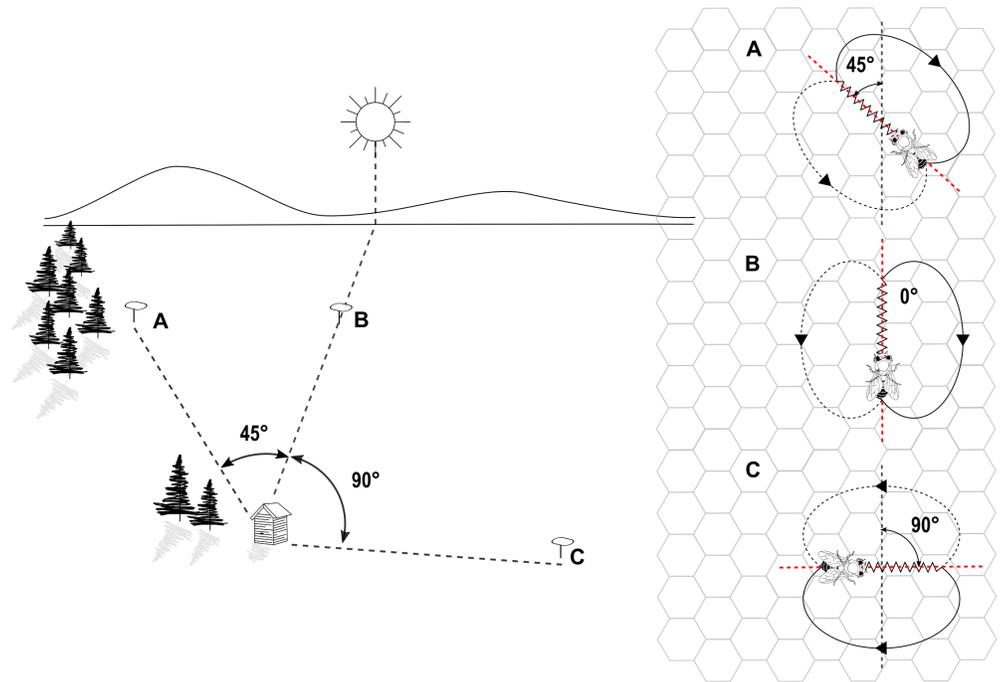


Figure 1. Correlation between waggle dance parameters and locations on the field. On the left, three food sources in the field located at A) 45° counterclockwise B) 0° and C) 90° clockwise, with respect to the azimuth. On the right, their representation through waggle dance paths on the surface of a vertical honeycomb.

Unemployed foragers might become interested in a waggle dance, follow the dancer's movements and decode the information contained in the dance to later search for the resource in the field [1, 3, 28, 32, 35]. Successful followers, once back in the hive, may also advertise the visited food source through dances, thereby incrementing the probability for their food source to recruit new foragers.

Study of the waggle dance as an abstract form of communication became a topic of great interest for ethologists after it was first described by von Frisch [43]. Keeping bee colonies in observation hives for their study is straightforward, and the complex dance behavior allows insights into many aspects of the honey bee biology. Still, several research fields focus on the bee dance and its encoded information.

The dance essentially contains an ordered pair of polar coordinates for a field location. Hence, waggle dances can be mapped back to the field with the help of pre-recorded calibration curves [42]. To obtain these curves, bees first are trained to example locations, then, waggle orientation and duration of their dances is noted. Once these functions have been obtained, the encoded location can be mapped into the field using the solar azimuth as a reference. Thus, without tracking the foragers' flight, one can deduce the distribution of foragers in the environment by establishing the distribution of dance-communicated locations. Couvillon et al. [7] used this method to investigate how the decline in flower-rich areas affects honey bee foraging, while Balfour and Ratnieks [2] used it to find new opportunities for maximizing pollination of managed honey bee colonies. Others have studied the information content of the dance

[19, 33] and the accuracy and precision with which bees represent spatial coordinates through waggle runs [8, 10]. Landgraf and co-workers tracked honey bee dances in video recordings to build a motion model for a dancing honey bee robot [25, 27]. Studies on honey bee collective foraging also focus on the waggle dance [34], including successive studies that model their collective foraging [5, 11] and nest-site selection behavior [30, 31, 36]. In [45] a previous version of the solution here described was used to automatically decode waggle dances as part of an integrated solution for the automatic long-term tracking of activity inside the hive.

Different techniques have been used over time to decode waggle dances. During the first decades that followed von Frisch’s discovery, most of the dances were analyzed in real time, directly from the observation hive with the help of protractors and stopwatches [35, 44]. Throughout the last decade, the use of digital video has become ubiquitous to extract the encoded information during a later process. Digital video allows researchers to analyze dances frame by frame and extract their characteristics either manually using the screen as a virtual observation hive [8], or assisted by computer software [10]. Although digital video recordings allow measurements with higher accuracy and precision, decoding communication dances continues to be a manual and time-consuming task.

Multiple automatic and semi-automatic solutions have been proposed to simplify and accelerate the dance decoding process. A first group of solutions focuses on mapping the bees’ trajectories via tracking algorithms [22, 23, 26]. A second group analyzes these trajectories and extracts specific features such as waggle run orientation and duration, using either a generic classifier trained on bee dances (see [14, 21, 29]) or methods based on hand-crafted features such as the specific spectral composition of the trajectory in a short window [27]. Although methods from the two groups have been combined in what is potentially an automatic detector and decoder of dances [15], its implementation has been limited to the automatic labeling of behaviors.

Here we propose a solution to detect and decode waggle dances automatically. Since all information known to be carried in the dances, can be inferred from the waggle run characteristics, our algorithm exclusively detects this portion of the dance directly from the video stream, avoiding a separate tracking stage. Our system detects 89.8% of all waggle runs with a false positive rate of only 5%. Compared to a human observer, the system extracts the waggle orientation with an average error of -2.92° ($SD = 7.37^\circ$) well within the range of human error.

Materials and Methods

Hive and recording setup

Our solution can be used either online with live streaming video or offline with recorded videos. The software requires a frame rate of approximately 100 Hz and a resolution of at least 1.5 pixels per millimeter. Thus QVGA resolution suffices to cover the whole surface of a “Deutschnormal” frame (370 mm x 210 mm). The four frame corners are used as a reference to rectify distortions caused by skewed viewing angles or camera rotations; If the frame corners are not captured by the camera, it is necessary to consider other reference points and their relative coordinates. The video setup used during the experiments is depicted in Fig 2 and might serve as a template for the interested researcher.

During our experiments, we used a one frame observation hive and one modified PS3eye camera per side. The camera is a low-cost model that offers frame rates up to 125 Hz at QVGA resolution (320 x 240 pixels) when using an alternative driver (see supplementary info S1 Text). For the lighting setup, it is necessary to use a constant

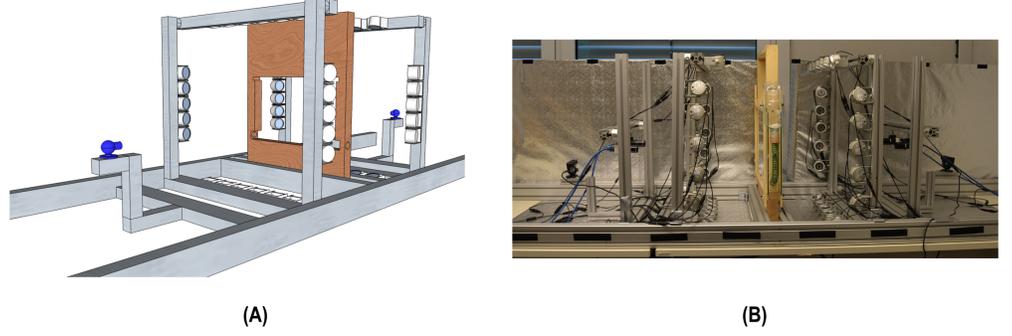


Figure 2. Recording setup. (A) A basic recording setup would consist of an observation hive, an array of LED clusters for illumination and one high-speed camera per side. (B) Our setup consisted of a one frame observation hive, an array of IR-LED clusters (840 nm wavelength) and one high-speed camera (a modified PS3Eye) per side. It was implemented alongside a second set of 4 cameras belonging to a different system. The observation hive was connected to one of the walls through a tunnel to allow the free transit of colony members between the hive and the outside.

light source, for instance, LEDs, another kind of sources such as fluorescent lamps may introduce flickering to the video, leading to suboptimal detection results. Our setup was illuminated by an array of infrared IR-LED clusters (840 nm wavelength). The entire structure was enveloped with a highly IR reflective foil with small embossments for light dispersion. The IR LED clusters pointed towards the foil to create a homogeneous ambient lighting and reduce reflections on the glass panes. The built in IR filter of the PS3eye cameras had to be removed to make them IR sensitive.

Target features

The relation of site properties: distance and direction to the feeder, and dance properties: duration and angle, have been recorded via systematic experiments [44]. These data can be used for an approximate inverse mapping of dance parameters to site properties.

$$r_R \sim f_d^{-1}(d_w). \quad (1)$$

$$\theta_R \sim \text{atan2} \left(\sum_{j=1}^n \sin \alpha_{wj}, \sum_{j=1}^n \cos \alpha_{wj} \right), n = 2k. \quad (2)$$

$$p_R \sim \frac{d_w}{d_r}. \quad (3)$$

Where r_R (Eq 1) is the distance between hive and resource and is related to the average waggle run duration d_w through the function f_d that approximates the calibration curve [42]. θ_R (Eq 2) is the angle between resource and solar azimuth (see Fig 1), and it corresponds to the average orientation of the waggle runs with respect to the vertical, with an even number of consecutive runs to avoid errors due to the divergence angle [27, 40, 44, 46]. The resource's profitability p_R (Eq 3) is proportional to the ratio between average waggle run duration d_w and average return run duration d_r . Dances for high-quality resources contain shorter return runs than those for less profitable resources located at the same distance, hence yielding a higher p_R value [37].

From Eq 1 to Eq 3 it follows that to decode the information contained in a communication dance three measurements are required: average waggle run duration d_w , average orientation α_w , and average return run duration d_r . In contrast to some approaches that require tracing the dance path to then analyze it and extract its characteristics [22, 23], we propose an algorithm that directly analyzes video frames to obtain each waggle run starting timestamp, duration, and angle. In our approach, return run durations are calculated as the time difference between the end of a waggle run and the beginning of the next one Fig 3.

Figure 3. Fundamental parameters. Knowing the starting time (t_x) and duration (d_{wx}) for each waggle run, it is possible to calculate the return run durations as the time gaps between consecutive waggle runs.

Software modules

Our software consists of four modules that are executed in sequence namely: attention module (AM), filter network (FN), waggle orientation module (OM) and mapping module (MM). The AM runs in real-time and stores small subregions of the video containing waggle-like activity. Later, false positives are filtered out using a convolutional neural network FN . The OM extracts the duration and angle of the waggle runs. Finally, the mapping module (MM) clusters waggle runs belonging to the same dances and maps their parameters back to field coordinates. All modules can run offline on video recordings. Long observations, however, require large storage space. Therefore, we propose using the AM with a real-time camera input to reduce the amount of stored data drastically. A detailed description of all four modules is given in the following sections.

Attention module (AM)

The waggle frequency of dancing bees lies within a particular range of values; we call the *waggle band* [17, 26]. Under consistent lighting conditions, the bee body is well discriminable from the background; therefore, brightness evolution of each pixel in the image originates from either honey bee activity or sensor noise. Thus, when a pixel intersects with a dancing bee, its intensity time-series is a function of the texture pattern on the bee and her motion dynamics. Indeed, by using a camera with low spatial resolution, bees appear as homogeneous ellipsoid blobs without surface texture. Thus the brightness of pixels that are crossed by wagging bees varies with the periodic waggle motion, and their frequency spectrum exhibits components in the *waggle band* or harmonics. To detect this general feature, we define a binary classifier here on referred to as *Dot Detector* (DD), each pixel position $[i, j]$ is associated to a DD D_{ij} . The DD analyzes the intensity evolution of the pixel within a sliding window of width b . For this purpose, the last b intensity values of each pixel are stored in a vector B , which at the time n can be described as $B_{ij}^n = [v_{ij}^{n-b+1}, v_{ij}^{n-b+2}, \dots, v_{ij}^n]$, where v_{ij}^k is the intensity value of the pixel $[i, j]$ at time k . We calculate a score for each of these time series using a number of sinusoidal basis functions, in principle similar to the Discrete Fourier Transform [6]:

$$score(\bar{B}_{ij}^n, r) = \sum_{m=1}^b \left(\left(\bar{B}_{ij}^n(m) \cdot \cos\left(2\pi r \frac{m}{s_r}\right) \right)^2 + \left(\bar{B}_{ij}^n(m) \cdot \sin\left(2\pi r \frac{m}{s_r}\right) \right)^2 \right) \quad (4)$$

where \bar{B}_{ij}^n is the normalized version of B_{ij}^n with $\bar{B}_{ij}^n \in \{-1, 1\}$ and $min_{ij}^n = min(B_{ij}^n)$ and $max_{ij}^n = max(B_{ij}^n)$, s_r is the video's sample rate (100 Hz), and

$r \in [10, 16]$ are the frequencies in the waggle band. If at least one of the frequencies in the *waggle band* scores over a defined threshold t_h , D_{ij} is set to 1. After computing the scores, those D_{ij} set to 1 are clustered together following a hierarchical agglomerative clustering (*HAC*) approach [38], using as a metric the Euclidean distance between pixels and with a threshold d_{max1} set to half the body length of a honey bee. Clusters formed by less than c_{min1} *DDs* are discarded as noise-induced, and the centroids of the remaining clusters are regarded as positions of potential dancers.

Positions found during the clustering step are then used to detect waggle runs (*WR*). If positions detected in successive frames are located within a maximum distance d_{max2} , defined according to the average waggle forward velocity (see [27]), the positions are considered as belonging to the same *WR*. At each iteration new dancer positions are matched against open *WR* candidates, and either appended to a candidate or used as basis for a new one. A *WR* candidate can remain open up to g_{max2} frames without new detections being added, after what it is closed. Only closed *WR* candidates with a minimum of c_{min2} detections are retrieved as *WRs*. Finally, coordinates of the potential dancer along with 50 x 50 pixels image snippets of the *WR* sequence are stored to disk.

The operation of the *AM* can be seen as a three layers process summarized in the following points:

1. Layer 0, for each new frame I^n :
 - (a) Update *DDs*' score vector.
 - (b) Set to 1 *DDs* with spectrum components in the waggle band above t_h .
2. Layer 1, detecting potential dancers:
 - (a) Cluster together *DDs* potentially activated by the same dancer.
 - (b) Filter out clusters with less than c_{min1} elements.
 - (c) Retrieve clusters' centroids as coordinates for potential dancers.
3. Layer 2, detecting waggle runs:
 - (a) Create waggle run assumptions by concatenating dancers positions with a maximum Euclidean distance of d_{max2} .
 - (b) Assumptions with a minimum of c_{min2} elements are considered as real *WR*.

Filtering with convolutional neural network (*FN*)

For long term observations we propose using the *AM* to filter relevant activity from a camera stream in real-time. This significantly reduces the disk space otherwise required to store full sized videos. However, depending on the task at hand, it may be advisable to chose the parameters of the module such that the rate of false negatives is minimal, yielding a number of false positive detections. A higher precision can be achieved by filtering the detections afterwards. To this end, we have trained a convolutional neural network that processes the sequence of 50 x 50 px images. The scalar output of the network is then thresholded to predict whether the input sequence contains a waggle dance. The network is a 3D convolutional network whose convolution and pooling layers are extended to the 3rd, i.e. temporal, dimension [20, 24, 39, 41]. The network architecture, three convolutional and two fully connected layers, is rather simple but suffices for the filtering tasks (for details refer to S2 Fig).

The network was trained on 8239 manually labeled *AM* detections from two separate days. During training, subsequences consisting of 128 frames were randomly sampled from the detections for each mini-batch. Detections with less than 128 frames were padded with constant zeros. Twenty percent of the manually labeled data was reserved

for validation. To reduce overfitting, the sequences were randomly flipped on the horizontal and vertical axes during training. We used the Adam optimizer to train the network and achieved an accuracy of 90.07% on the validation set. This corresponds to a recall of 89.8% at 95% precision.

Orientation module (*OM*)

While the duration of a *WR* is estimated from the number of frames exported by the *AM*, its orientation is computed in a separate processing step here on referred as orientation module (*OM*), usually performed offline to keep computing resources free for detecting waggle runs. Dancing bees move particularly fast during waggle runs, throwing their body from side to side at a frequency of about 13 Hz [27]. Images resulting from subtracting consecutive video frames of wagging bees exhibit a characteristic pattern similar to a 2D Gabor filter, a positive peak next to a negative peak, which orientation is aligned with the dancer’s body Fig 4.

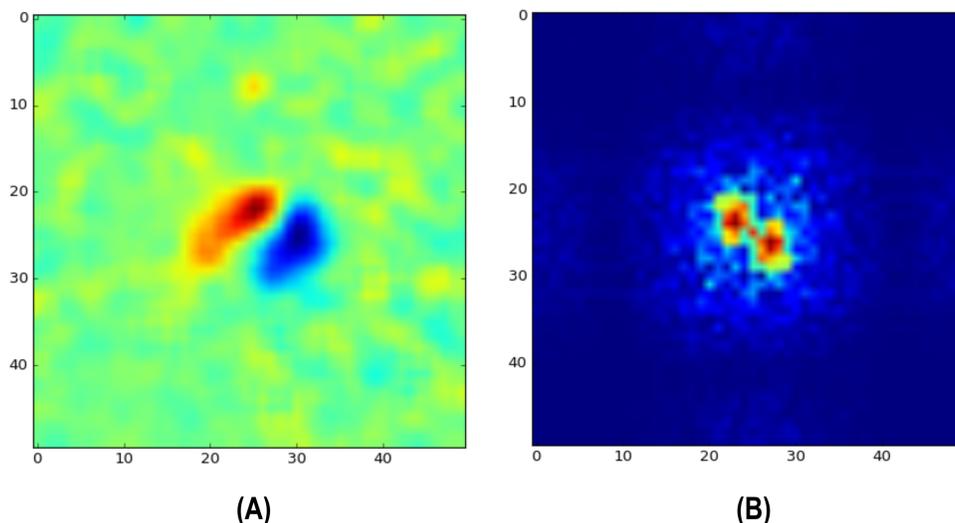


Figure 4. Difference image and its Fourier transformation. (A) The image resulting from subtracting consecutive video frames of wagging bees exhibits a characteristic Gabor filter-like pattern. (B) While the peak location varies in image space along with the dancer’s position, its representation in the Fourier space is location-independent, showing distinctive peaks at frequencies related to the size and distance of the Gabor-like pattern.

The Fourier transformation of the difference image provides a location-independent representation Fig 4B of the wagging event while preserving information regarding the dancer’s orientation. We make use of the Fourier slice theorem [4], which states that the Fourier transform of a projection of the original function onto a line at an angle α is just a slice through the Fourier transform at the same angle. Imagine a line orthogonal to the dancer’s orientation. If we project the Gabor-like pattern onto this line, we obtain a clear sinusoidal pattern which appears as a strong pair of maxima in the Fourier space at the same angle. Not all difference images in a given image sequence exhibit the Gabor pattern, it only appears when the bee is quickly moving laterally. To get a robust estimate of the waggle orientation, we sum all Fourier transformed difference images Fig 5A and apply a bandpass filter Fig 5B to obtain the correct maxima locations Fig 5C. The bandpass filter is performed in the frequency domain by multiplying with a difference-of-Gaussians *DoG*. The radius of the ring needs to be

tuned to the expected frequency of the sinusoidal in the 1D projection of the Gabor pattern. This frequency depends on the frame rate (in our case 100 Hz) and the image resolution (17 px/mm). With the lateral velocity of the bee (we used the descriptive statistics in [27]) one can compute the displacement in pixels (5 to 7 px/frame). Using Eq 5 we can approximate the value of the expected frequency k :

$$k = \frac{I_{size}}{T} = \frac{I_{size}}{2x}, \quad (5)$$

where I_{size} is the input image size (50 px in this case) and T is the period for the Gabor filter-like pattern or twice the bee’s displacement between frames.

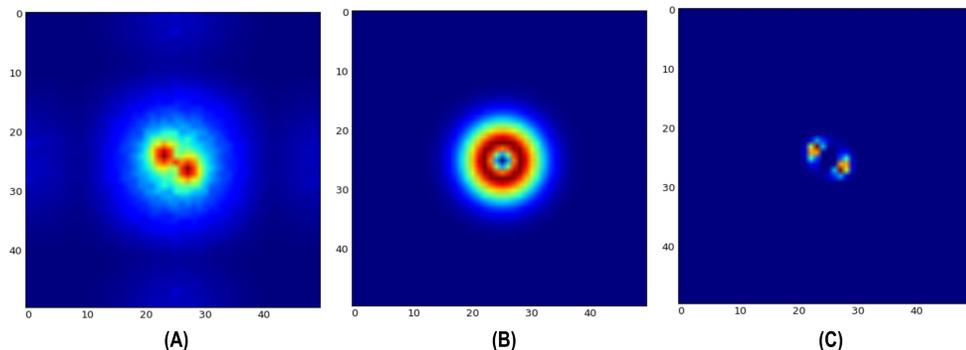


Figure 5. Filtering cumulative sum of difference images in the Fourier space. (A) The cumulative sum of the Fourier transformed difference images of a waggle run exhibit a strong pair of maxima in locations orthogonal to the dancer’s orientation. (B) A *DoG* kernel of the Mexican hat type, properly adjusted to the waggle band, is used as a bandpass filter. (C) Bandpass filtering the cumulative sums emphasizes values within the frequencies of interest.

The orientation of the waggle run is obtained from the resulting image through Principal Component Analysis (*PCA*). However, the principal direction reflects the direction of the dancer’s lateral movements, so it is necessary to add 90° to this direction to obtain the dancer’s body axis. This axis represents two possible waggle directions. To disambiguate the alternatives, we process the dot detector positions extracted by the *AM*. Each of these image positions represents the average pixel position in which we found brightness changes in the waggle band. In a typical waggle sequence, these points trace roughly the path of the dancer. We average all *DD* positions of the first 10% of the waggle sequence and compute all *DD* positions relative to this average. We then search for the maximum values in the histogram of the orientation of all vectors and average their direction for a robust estimate of the main direction of the dot detector sequence. This direction is then used to disambiguate the two possible directions extracted by *PCA*.

Mapping module (*MM*)

Waggle dances encode polar coordinates for field locations. To map these coordinates back to the field we implemented a series of steps in what is here on referred as mapping module (*MM*). The *MM* reads the output of the *AM* and *OM*, essentially time, location, duration and orientation of each detected waggle run. Then, waggle runs are clustered following a *HAC* approach, similar to the *AM*. In this case, the clustering process is carried out in a three-dimensional data space defined by the axes X and Y of the comb surface and a third axis T of time of occurrence. This way, each *WR* can be represented in the data space by (x, y, t) coordinates based on its comb location and

time of occurrence (see S3 FigA). To maintain coherence between spatial and temporal values, the time of occurrence is represented in one fourth of the seconds relative to the beginning of the day.

A threshold Euclidean distance d_{max3} is defined as a parameter for the clustering process (see S3 FigB). The value of the threshold is based on the average drift between *WRs* and the average time gap between consecutive *WRs* (we used the data provided in [27]). Only clusters with a minimum of 4 waggle runs are considered as actual dances. Then, we use random sample consensus (*RANSAC*) [16] to find outliers in the distribution of waggle run orientations. Waggle run duration and orientation are then averaged for all inliers and translated to field locations. The mean waggle run duration is translated to meters using a calibration curve, and its orientation is translated to the field with reference to the azimuth at the time of the dance.

Experimental validation and results

To evaluate the performance of the *OM*, we compared the results obtained for a set of 200 waggle runs against manually defined angles measured by eight human observers. A custom user interface allowed the users to trace the line that best fits the dancer’s body (see supplementary info S2 Text). The reference angle for each waggle run was defined as the average of the eight manually extracted angles. The *OM* performed with an average error of -2.92° and an SD of 7.37° , close to the SD of 6.66° observed in the human generated data (further details in supplementary info S3 Text).

To illustrate the use cases of the automatic detection of waggle dances we mapped all dances detected by our system during a period of 5 hours Fig 6. The data was collected from a honey bee colony kept under constant observation during the summer of 2016. A group of foragers from the colony was trained to an artificial feeder placed 325 m southwest of the hive. We used dance recordings from a previous experiment in the same area to calibrate the distance/waggle run duration parameter [27]. On the test day, a total of 571 dances were detected with an average of 5.8 waggle runs per dance.

Since colony foragers were allowed to forage from other food sources, not all of the detected dances point towards the artificial feeder. However, most of the detected dances cluster around the feeding site. Averaging the direction of all detected dances we obtained a very precise match with the artificial feeder’s direction, with a negligible error below 0.17° . The spread of points around the feeder matches the expected angular standard deviation [27].

Discussion

We have presented the first automatic waggle dance detection and decoding system. It is open source and available for free. It does not require expensive camera hardware and works with standard desktop computers. The system can be used continuously for months, and its accuracy is close to human performance ($M = -3.3^\circ$, $SD = 5.5^\circ$ vs. $M = 0^\circ$, $SD = 3.7^\circ$). A dance deviating 3° from the correct direction is mapped with a 15 m error at 300 m distance, growing linearly with distance. On the waggle run level, the standard error is larger due to waggle orientations decoded with an 180° error. We found that the orientation of 9% of the waggle runs in our test set are incorrectly flipped. This flip error is corrected by clustering waggle runs to dances and removing angular outliers with *RANSAC*. Once the outliers have been removed, our system performs with an average error of -2.0° ($SD = 6.1^\circ$) on the waggle run level.

Our dataset exhibits dances for whose waggle runs the *OM* returns many angles flipped. These dances were automatically removed since no strong mode could be

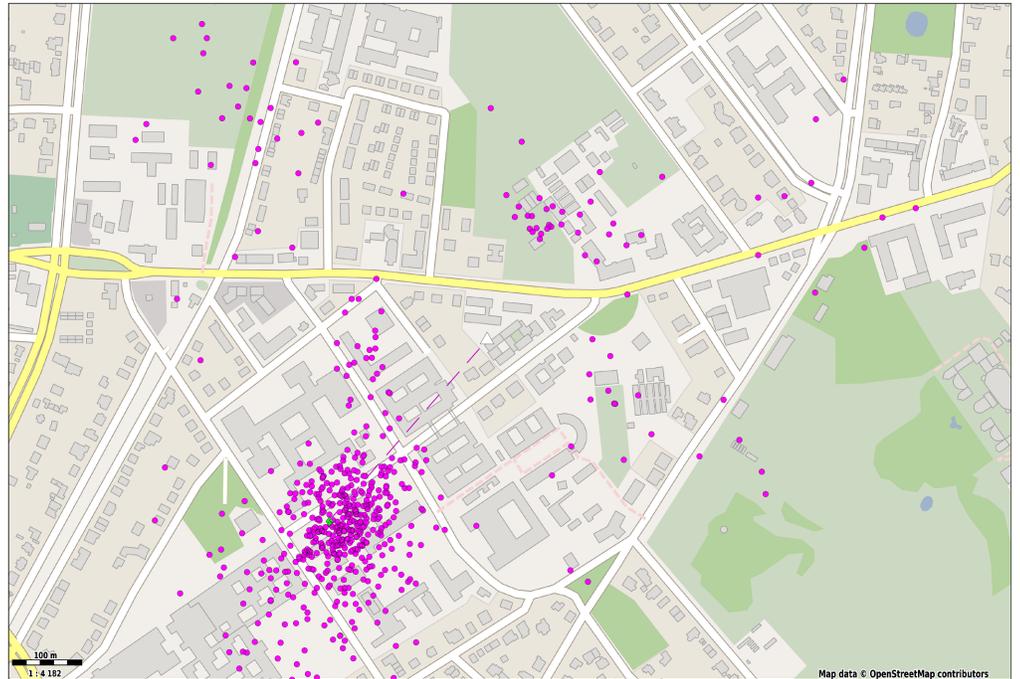


Figure 6. Detected dances mapped back to the field. Dance detections were averaged over at least four waggle runs and translated to a field location with respect to the sun’s Azimuth. A linear mapping was used to convert waggle duration to metric distance. The hive and feeder positions are depicted with a white and green triangle, respectively. The dashed line represents the average dance direction.

identified. The angular error on the waggle run level is also larger as when manually removing the flipped waggle runs (SD: 7.37°). Thus, the extraction of the correct angle can be further improved by perfecting the method for disambiguating between the reported angle and the flipped alternative. We inspected waggle runs that were incorrectly decoded and found that in most of these examples the dancing bee is partly occluded or waggles for short durations only, i.e. there is almost no forward motion visible. The forward motion is the central feature in the orientation reader module used to disambiguate the direction. Recognizing anatomical features of the dancing bee, such as the head or abdomen, could help reduce this common error. With falling costs for better camera and computing hardware in mind, we, however, think that just using a higher spatial resolution will likely resolve most of the detection and decoding errors we have described.

The proposed system consists of multiple modules which might seem impractical or even obfuscating to the end-user. We have documented our software carefully and believe this system will be of great help especially when decoding of hundreds of dances is required. The modules are command line programs. We feel this might reduce acceptance with biologist and are thus developing a graphical user interface to be published in the near future. We currently investigate whether a deep convolutional network is able to extract the relevant image features. If successful, this would enable us to merge the filter network module and the orientation reader, therefore reducing complexity for the user. For the future, we envision an entirely neural system for all the described stages. We also think the solution could be ported to mobile devices. This would enable users an easier setup. Dance orientations could be corrected by reading the direction of gravity directly from the built-in accelerometer. We would like to

encourage biologists to use our system and report issues that they face in experiments.

Supporting Information

Optionally make clicky links to the online file using `\href{url}{description}`.

S1 Text. Specifications of the recording setup used during reported experiments. This document contains further information on the recording setup, with an emphasis on technical details.

S2 Text. Further details on the software modules. This document contains diagrams and detailed information on the functioning of the software modules.

S3 Text. Error distributions at multiple levels of analysis. This document provides additional results supporting the case of study presented in the experimental validation and results section.

S2 Fig. A convolutional neural network was used to filter the detections of the AM. The raw sequences of images are processed by two stacked 3D convolution layers with SELU nonlinearities. The outputs of the second convolutional layer are flattened using average pooling on all three dimensions. A final fully connected layer with a sigmoid nonlinearity computes the probability of the sequence being a dance or not. Dropout is applied after the average pooling operation to reduce overfitting.

S3 Fig. Representation of WRs in the data space XYT. (A) Representation of a set of 200 WRs in the XYT data space, where values in the axes X and Y are defined by their comb location, and in axis T by their time of occurrence. (B) WRs within a maximum Euclidean distance of d_{max3} are clustered together and regarded as dances.

Acknowledgments

The authors thank M. Halilovic, F. Lojewski, A. Rau and S. Witte for their contribution and support conducting the experiments and in the development of this system. This work is funded in part by the German Academic Exchange Service (DAAD).

References

1. H. Al Toufailya, M. Couvillon, F. L. W. Ratnieks, and C. Grüter. Honey bee waggle dance communication: Signal meaning and signal noise affect dance follower behaviour. *Behavioral Ecology and Sociobiology*, 67(4):549–556, 2013.
2. N. J. Balfour and F. L. W. Ratnieks. Using the waggle dance to determine the spatial ecology of honey bees during commercial crop pollination. *Agricultural and Forest Entomology*, dec 2016.
3. J. C. Biesmeijer and T. D. Seeley. The use of waggle dance information by honey bees throughout their foraging careers. *Behavioral Ecology and Sociobiology*, 59:133–142, 2005.

-
4. R. Bracewell. The Projection-Slice Theorem. In *Fourier Analysis and Imaging*, pages 493–504. Springer US, Boston, MA, 2003.
 5. S. Camazine and J. Sneyd. A model of collective nectar source selection by honey bees: Self-organization through simple rules. *Journal of Theoretical Biology*, 149(4):547–571, 1991.
 6. J. W. Cooley and J. W. Tukey. An Algorithm for the Machine Calculation of Complex Fourier Series. *Mathematics of Computation*, 19(90):297, apr 1965.
 7. M. Couvillon, R. Schürch, and F. L. W. Ratnieks. Waggle Dance Distances as Integrative Indicators of Seasonal Foraging Challenges. *PLoS ONE*, 9(4):e93495, 2014.
 8. M. J. Couvillon, F. C. Riddell Pearce, E. L. Harris-Jones, A. M. Kuepfer, S. J. Mackenzie-Smith, L. A. Rozario, R. Schürch, and F. L. W. Ratnieks. Intra-dance variation among waggle runs and the design of efficient protocols for honey bee dance decoding. *Biology open*, 1(5):467–72, may 2012.
 9. M. Dacke and M. V. Srinivasan. Two odometers in honeybees? *Journal of Experimental Biology*, 211(20):3281–3286, 2008.
 10. R. J. De Marco, J. M. Gurevitz, and R. Menzel. Variability in the encoding of spatial information by dancing bees. *The Journal of experimental biology*, 211(Pt 10):1635–1644, may 2008.
 11. H. de Vries and J. C. Biesmeijer. Modelling collective foraging by means of individual behaviour rules in honey-bees. *Behavioral Ecology and Sociobiology*, 44(2):109–124, nov 1998.
 12. H. E. Esch and J. E. Burns. Distance estimation by foraging honeybees. *The Journal of experimental biology*, 199:155–62, 1996.
 13. H. E. Esch, S. Zhang, M. V. Srinivasan, and J. Tautz. Honeybee dances communicate distances measured by optic flow. *Nature*, 411(6837):581–583, may 2001.
 14. A. Feldman and T. Balch. Automatic Identification of Bee Movement. In *International Workshop on the Mathematics and algorithms of social insects*, pages 53–59, Georgia, 2003. Georgia Institute of Technology.
 15. A. Feldman and T. Balch. Representing Honey Bee Behavior for Recognition Using Human Trainable Models. *Adaptive Behavior*, 12(3-4):241–250, dec 2004.
 16. M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, jun 1981.
 17. M. Gil and R. J. De Marco. Decoding information in the honeybee dance: revisiting the tactile hypothesis. *Animal Behaviour*, 80:887–894, 2010.
 18. C. Grüter, M. S. Balbuena, and W. M. Farina. Informational conflicts created by the waggle dance. *Proceedings. Biological sciences / The Royal Society*, 275(March):1321–1327, 2008.
 19. J. B. S. Haldane and H. Spurway. A statistical analysis of communication in “*Apis mellifera*” and a comparison with communication in other animals. *Insectes Sociaux*, 1(3):247–283, sep 1954.

-
20. S. Ji, W. Xu, M. Yang, and K. Yu. 3D Convolutional Neural Networks for Human Action Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):221–231, jan 2013.
 21. M. Kabra, A. A. Robie, M. Rivera-Alba, S. Branson, and K. Branson. JAABA: interactive machine learning for automatic annotation of animal behavior. *Nature Methods*, 10(1):64–67, 2012.
 22. Z. Khan, T. Balch, and F. Dellaert. A Rao-Blackwellized Particle Filter for Eigentracking. In *2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 980–986. IEEE, 2004.
 23. T. Kimura, M. Ohashi, R. Okada, and H. Ikeno. A new approach for the simultaneous tracking of multiple Honeybees for analysis of hive behavior. *Apidologie*, 42(5):607–617, 2011.
 24. G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter. Self-Normalizing Neural Networks. *arXiv preprint*, jun 2017.
 25. T. Landgraf, M. Oertel, A. Kirbach, R. Menzel, and R. Rojas. Imitation of the honeybee dance communication system by means of a biomimetic robot. In *Lecture Notes in Computer Science*, volume 7375 LNAI, pages 132–143. Springer-Verlag, 2012.
 26. T. Landgraf and R. Rojas. Tracking honey bee dances from sparse optical flow fields. Technical Report June, Freie Universität Berlin, Berlin, 2007.
 27. T. Landgraf, R. Rojas, H. Nguyen, F. Kriegel, and K. Stettin. Analysis of the waggle dance motion of honeybees for the design of a biomimetic honeybee robot. *PLoS ONE*, 6(8):e21354, 2011.
 28. R. Menzel, A. Kirbach, W. D. Haass, B. Fischer, J. Fuchs, M. Koblofsky, K. Lehmann, L. Reiter, H. Meyer, H. Nguyen, S. Jones, P. Norton, and U. Greggers. A common frame of reference for learned and communicated vectors in honeybee navigation. *Current Biology*, 21(8):645–650, 2011.
 29. S. M. Oh, J. M. Rehg, T. Balch, and F. Dellaert. Learning and Inferring Motion Patterns using Parametric Segmental Switching Linear Dynamic Systems. *International Journal of Computer Vision*, 77(1-3):103–124, may 2008.
 30. K. M. Passino and T. D. Seeley. Modeling and analysis of nest-site selection by honeybee swarms: the speed and accuracy trade-off. *Behav Ecol Sociobiol*, 59:427–442, 2006.
 31. A. Reina, G. Valentini, C. Fernández-Oto, M. Dorigo, and V. Trianni. A Design Pattern for Decentralised Decision Making. *PLoS ONE*, 10(10):e0140950, oct 2015.
 32. J. R. Riley, U. Greggers, a. D. Smith, D. R. Reynolds, and R. Menzel. The flight paths of honeybees recruited by the waggle dance. *Nature*, 435(7039):205–207, 2005.
 33. R. Schürch and F. L. W. Ratnieks. The spatial information content of the honey bee waggle dance. *Frontiers in Ecology and Evolution*, 3:22, mar 2015.
 34. T. D. Seeley. Social foraging by honeybees: how colonies allocate foragers among patches of flowers. *Behavioral Ecology and Sociobiology*, 19(5):343–354, nov 1986.

-
35. T. D. Seeley. *The Wisdom of the Hive*. Harvard University Press, 1995.
 36. T. D. Seeley and S. C. Buhrman. Nest-site selection in honey bees: How well do swarms implement the "best-of-N" decision rule? *Behavioral Ecology and Sociobiology*, 49:416–427, 2001.
 37. T. D. Seeley, A. S. Mikheyev, and G. J. Pagano. Dancing bees tune both duration and rate of waggle-run production in relation to nectar-source profitability. *Journal of Comparative Physiology - A Sensory, Neural, and Behavioral Physiology*, 186(9):813–819, oct 2000.
 38. R. Sibson. SLINK: An optimally efficient algorithm for the single-link cluster method. *The Computer Journal*, 16(1):30–34, jan 1973.
 39. N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.
 40. D. A. Tanner and P. K. Visscher. Adaptation or constraint? Reference-dependent scatter in honey bee dances. *Behavioral Ecology and Sociobiology*, 64(7):1081–1086, 2010.
 41. D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning Spatiotemporal Features with 3D Convolutional Networks. *arXiv preprint*, dec 2014.
 42. P. K. Visscher and T. D. Seeley. Foraging Strategy of Honeybee Colonies in a Temperate Deciduous Forest. *Ecology*, 63(6)(6):1790–1801, 1982.
 43. K. von Frisch. Die Tänze der Bienen. *Österreich Zool . Z.*, 1:1–48, 1946.
 44. K. von Frisch. *Tanzsprache und Orientierung der Bienen*. Springer, Berlin, 1965.
 45. F. Wario, B. Wild, M. Couvillon, R. Rojas, and T. Landgraf. Automatic methods for long-term tracking and the detection and decoding of communication dances in honeybees. *Frontiers in Ecology and Evolution*, 3(September):1–14, 2015.
 46. A. Weidenmüller and T. D. Seeley. Imprecision in waggle dances of the honeybee (*Apis mellifera*) for nearby food sources: Error or adaptation? *Behavioral Ecology and Sociobiology*, 46(3):190–199, 1999.