Freie Universität Berlin

Master's Thesis at the Department for Informatics and Mathematics

Workgroup Artificial and Collective Intelligence

Correlation of a bee's visual perception with its brain activity

Tobias Schülke
Matriculation Number: 4767145
tobias.schuelke@fu-berlin.de

Supervisor & First Examiner: Prof. Dr. Tim Landgraf
Second Examiner: Prof. Dr. Dr. h.c. habil. Raúl Rojas

Berlin, October 28, 2019

Abstract

Despite their small brain size, honeybees are able to orientate themselves exceptionally well in the environment. In order to study their navigational capabilities, correlations between brain activity of a bee and its visual perception are analysed. Neural activity of the animal was recorded while flying with a quadcopter. First, it is shown that similar brain activity was measured in repeated flight routes. Next, the flight of the copter is simulated in a virtual 3D environment of the experiment's fields and its surroundings. Images were rendered from the perspective of a bee and processed by an autoencoder. Activations of the encoder were correlated with the bee's neural activity. Additionally, the brain activity was projected onto a field in the range of a bee's field of view. Both approaches revealed high correlations during flight turns.

Eidesstattliche Erklärung

Ich versichere hiermit an Eides Statt, dass diese Arbeit von niemand anderem als meiner Person verfasst worden ist. Alle verwendeten Hilfsmittel wie Berichte, Bücher, Internetseiten oder ähnliches sind im Literaturverzeichnis angegeben, Zitate aus fremden Arbeiten sind als solche kenntlich gemacht. Die Arbeit wurde bisher in gleicher oder ähnlicher Form keiner anderen Prüfungskommission vorgelegt und auch nicht veröffentlicht.

October 28, 2019

Tobias Schülke

Contents

1 Introduction

Honeybees, *Apis mellifera*, are well-known for their navigational capabilities. After visiting different food sources, they fly back to their hive and later find these locations again. Distances between such locations can be up to ten kilometres [1]. That is, the animals travel to desired destinations which are not in their viewing range. How bees are able to achieve this is still under research.

With a volume of about $1\,mm^3$ and fewer than a million neurons bees have a small brain compared to other animals which are known for orientating themselves in the environment like mice. Larger brains contain more replicated neuronal circuits whereby they, amongst others, have a greater sensory precision. But that does not necessarily lead to a higher intelligence [8]. Hence, the simpler brain of a bee is studied in which it might be easier to find structures related to navigation.

Young bees explore the environment in short-range flights around the hive and in long-range flights. Multiple successive orientation flights are performed to investigate different parts of the terrain surrounding the hive [11]. Learned characteristics of the hive or terrain during the first orientation flight are sufficient for bees to reach their home successfully after being displaced [12]. Similarly, re-orientation flights help bees to receive hive and landscape features when the hive is moved. Even in a new area one re-orientation flight is enough to find the way home from a random location [10].

There are multiple publications which indicate that bees use a cognitive map [5, 24, 25]. The idea is that bees are able to remember spatial relations between landmarks such that they locate themselves and head towards a destination. This theory was disputed [7] but there is further evidence which supports a cognitive map [4]. Landmarks are objects which are salient, permanent, and relevant such that navigational agents remember and recognize them, for example by their colour, size, edge orientation, and symmetry [16]. When bees fly home, they travel along elongated ground structures [26]. Moreover, it was shown that honeybees are able to fly novel shortcuts between known locations they visited separately by using path integration [25].

Honeybees inform foragers about their findings by performing a waggle-dance which encodes a direction vector and distance to flowers. Bees use the position of the sun during dances to describe the direction of food sources. This works even without direct visual contact between bee and sky [22]. The animals do not communicate the distance to a food source in absolute values during their dances but via the measured optical flow, which is the image motion during a flight, from hive to destination [14].

The workgroup *Artificial and Collective Intelligence* as well as the Institute of Neurobiology at Freie Universität Berlin research together on the project NeuroCopter. They examine the navigation skills of honeybees. The aim of this project is to understand how different parts of a bee's brain process and represent the animal's navigational capabilities. Likewise, their research could reveal basic computational structures which can be used for general navigation tasks [20].

Initially, the idea of the project was to imitate a bee with a custom built quadcopter, named NeuroCopter. In order to be able to fly autonomously, the logic of the copter consisted of multiple modules with artificial neural networks which represented the bee's brain. These modules were tested in computer simulations and trained with virtual sensory data. Afterwards, they were deployed to the copter. In a following approach, a purchased quadcopter was modified in such a way that a living honeybee is attached to it. An on-board computer on the copter records the brain activity of the bee during flights.

Experiments with NeuroCopter are analysed in this thesis. The goal was to find locations which the animal uses to navigate. It is shown that neural signals of the same bee are similar during iterations of the same experiment. Specific locations of the environment which trigger brain activity could not be found.

In the following section an overview is given of related work to navigation of honeybees and previous works of the NeuroCopter project. Next, Section 3 describes the experiment with NeuroCopter, how the bee's brain activity is processed for analyses in this thesis, as well as the environment of the experiment. The brain activity of the bee is examined in Section 4. Section 5 and Section 6 describe two approaches to correlate neural activity with visual features of the environment. Finally, the results are evaluated in Section 7 and a conclusion is drawn in Section 8.

2 Related Work

The visual perception and navigational capabilities of honeybees have been studied for several decades with different approaches. Menzel et al. used a harmonic radar system to track the flight paths of bees on a field with a hive, coloured tents used as landmarks, and a feeding station [24]. By varying the position of the hive, feeder, tents, and radar station as well as grouping bees with different knowledge of the environment, similarities between flights in multiple situations were compared. Straight flights of a learned direction from another bee, search flights, and goal-driven flights could be distinguished. Moreover, a radar system was used to show that bees use elongated ground structures for navigation [26].

The visual perception of honeybees can be investigated with tunnels. Chen et al. used a maze shaped like a Y to test the pattern recognition of bees [6]. Bees were shown two different patterns at both ends of the Y-maze while only one ending rewarded them with sugar water. After a training period different patterns were displayed. The bees were able to differentiate between topologically different patterns during training and abstract them to novel patterns. In another experiment, Si et al. showed that the bee's visually driven measurement which they need for their waggle dance is robust [38]. Although flying through a tunnel with varying and sparse texture, bees could use optic flow for distance measurement. Buatois et al. conducted two experiments: bees walking in a Y-maze and on a treadmill while in both scenarios visual stimuli where presented [3]. In the maze experiment the bees' visual learning performance was better. They propose that the mobility is necessary to learn the stimuli from different viewing angles.

Multiple studies focus on the change of brain activity while bees navigate. Paulk et al. measured neurons in different regions of a bumblebee's brain while the animal was presented an array of light emitting diodes (LED) to trigger colour stimuli as well as a cathode ray tube (CRT) monitor for motion stimuli [28]. Along the visual pathway from eye to higher order brain centres colour sensitivity is processed by neurons before motion stimuli. Similarly, in two experiments by Ibbotson et al. CRT displays were placed in front and on both sides of a honeybee [19]. High contrast sine-wave gratings or spiral patterns on the screens were shown for visual stimulation. Measured neurons during recordings with both screens on side suited to the optic flow. Mertes et al. recorded learning flights of bumblebees with two cameras in a laboratory arena with a custom-made texture on walls, a feeder and three cylinders which served as landmarks [27]. These recordings were used to obtain a three-dimensional flight trajectory. Together with a virtual model of the arena, the flight could be reconstructed in ego-perspective of the bumblebee.

This video was shown on a LED screen to a tethered bumblebee while recording neural activity in its brain. The virtual environment was modified by removing landmarks and changing texture of objects or the background. These modifications were presented to bumblebees, too. It turned out that the recorded neurons process landmarks. On the other hand, changes of textures did not have a great impact.

In the beginning, the goal of the NeuroCopter project was to develop a quadcopter which behaves like a honeybee. Autonomous flights were developed by connecting the onboard computer of the copter to a virtual flight simulation [9]. The simulation prevented hardware damage during flight tests. The multicopter was constructed individually [23]. It consisted of the on-board computer ODROID-U3 as well as multiple sensors such as a camera and a gyroscope. While flying the ODROID-U3 sent recorded data to a computer on the ground.

Various works focused on the implementation of algorithms which improved the behaviour of the copter. Using a neural network, the copter's ego motion could be extracted successfully from sparse optical flow fields [20, 40]. That means, the input only consists of camera images and no further sensors of the copter. The authors were able to recreate the flight curve with the help of the ego motion calculated by the network. This approach provides basis for navigation models of the copter since they need a way to determine the copter's ego motion. Another work focused on the determination of the relative position of the copter [2]. This was achieved by evaluating the copter's sensors and detection of landmarks on camera images. Furthermore, collision detection was developed [41]. The system recognizes dangerous situations by analysing the optical flow. It compares the situation to previously supervised learned data.

Lastly, the goal of the project was adjusted. A quadcopter was customized in such a way that a living honeybee can be attached to it [29]. During flights the copter measures neural activity of the bee's brain. With this system the navigation capabilities of flying bees can be examined in their natural environment. This was not possible with previously published experiments such as a simulation of the environment with screens and airflow machines to stimulate the bee.

3 Data Collection

Recording brain activity of a flying honeybee over long distances is difficult due to its small size. Julian Petrasch modified a multicopter during his master thesis at Biorobotics lab to solve this task [29]. Moreover, he flew the copter with a bee in field experiments. The captured brain activity is used in this thesis. This section summarizes aspects of the master thesis of Julian Petrasch which form the basis for the analysis in the next sections.

3.1 Multicopter with honeybee

A multicopter is well suited for measuring neuronal activity of a bee during a flight. It can fly as fast as bees and counterbalance wind. In addition, the copter is able to fly a programmed route autonomously such that experiments can be planned accurately and are repeatable. In this project a *DJI Matrice 100*[1] was selected which is a ready to fly copter for developers. This drone was modified as shown in Figure 1. The bee is attached to a gimbal on a cantilever. Because of that the animal is not distracted by the propellers of the copter. The gimbal keeps the bee in its natural position while flying since the copter tilts differently than bees do.

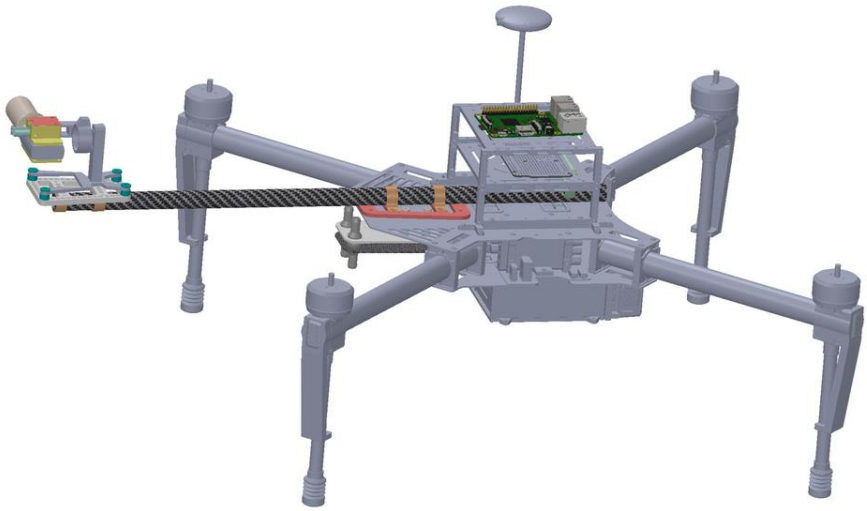

Figure 1: CAD model of the modified multicopter. A gimbal at the end of a cantilever holds a honeybee and aligns it in its natural position while flying. (Image from [29])

[1] Dji matrice 100. SZ DJI Technology Co., Ltd., 2018 (https://www.dji.com/de/matrice100)

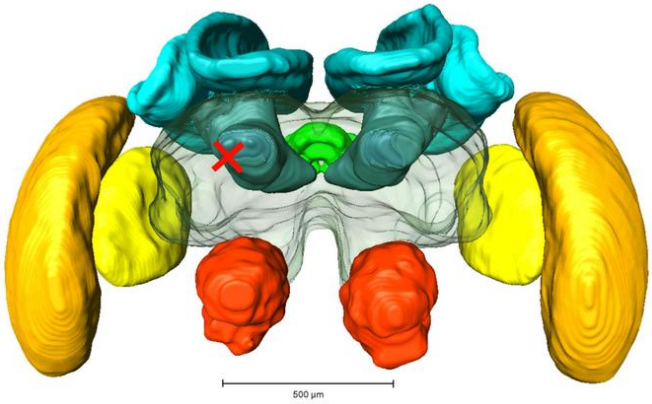

Figure 2: The symmetric brain of a bee consists of different parts: *Red, Antennal lobes, AL; Blue, Mushroom bodies, MB (comprising medial and lateral calyx and the pedunculus); Green, central complex, CB (comprising upper and lower central complex and protocerebral bridge); Yellow, Lobula; Orange, Medulla.* The Lobula and Medulla regions of the brain transfer visual signals to the mushroom body [30]. The electrode is placed at an exit of the mushroom body, marked with a red cross. (Description and image from [39], electrode position from [29])

In order to fix a honeybee to the copter, first a bee is cooled down for ten minutes to prevent it from moving. Then, the animal is put in a holder on the copter, the head is opened and an electrode is moved into the brain. The electrode is glued at an exit of a mushroom body, see Figure 2. This region of the brain receives direct visual input and is responsible for learning of visual information [30]. The bee is able to see the environment but it cannot move as depicted in Figure 3.

The analogue signal of the electrode in the brain of the honeybee is amplified by two amplifier boards. Next, the signal is digitized by analogue to digital converters and saved to log files on a SD-Card by a microcontroller. Grounding and shielding, for example with copper foil, is used to eliminate noise generated by the motors of the copter in the signal as much as possible. Besides capturing the neural activity of the honeybee, the drone records telemetry data such as its GPS position, orientation and speed. This is written to its internal memory.

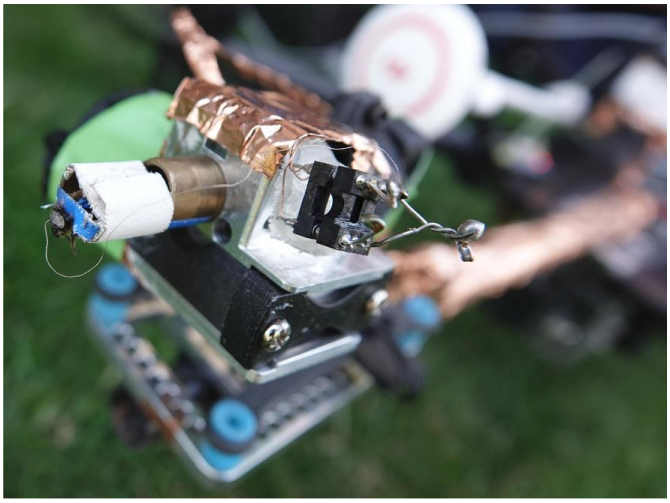

Figure 3: Honeybee is fixed at the gimbal on the multicopter. The bee's head is visible and connected to an amplifier board via a cable. (Image from [29])

3.2 Data processing

Sensor data of the drone as well as brain activity of the bee are stored as binary files on the internal memory and SD-Card. These information need to be converted into a readable format and combined in a time-synchronized manner. The telemetry data is converted to a CSV-file with *DatCon*[2] and synchronised with the brain activity with *NeuroCopter Binary Converter* [29] which was written for this project by Julian Petrasch. The software *Spike2*[3] is used to sort and classify the neural spikes of the combined CSV-file. Lastly, the detected spikes and their corresponding timestamps are again merged with the CSV-file containing the drone and brain data with *NeuroCopter Binary Converter*. The resulting CSV-file is the basis for the analysis in this thesis.

3.3 Field experiment

In 2018, experiments with the NeuroCopter were performed at Julius Kühn-Institut in Berlin by Julian Petrasch [29]. A food source was placed on a field at the institute. Bees that found the feeder were marked and later used for the NeuroCopter. After taking off and quickly flying to the start point, the flight route started from the same direction to the feeder as the bees flew previously from the hive, see Figure 4a. During the flight the copter takes three half circles and flies three times over the feeder from different directions.

[2] https://datfile.net/DatCon/intro.html
[3] http://ced.co.uk/products/spkovin

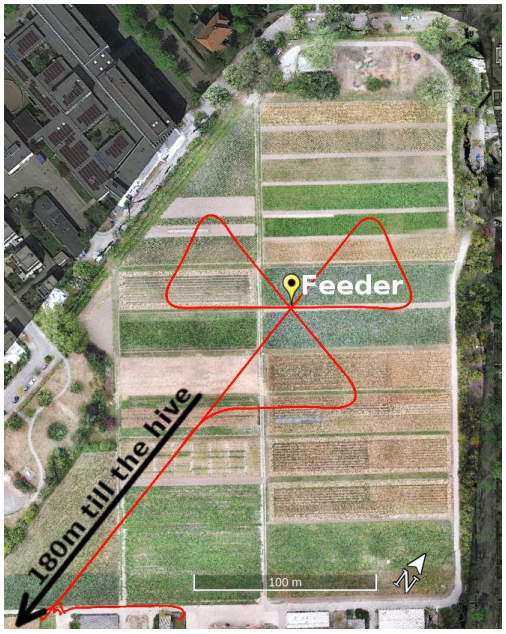

(a) 2nd August, 2018

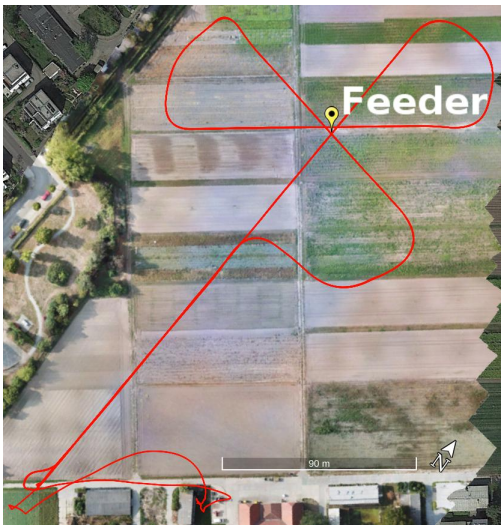

(b) 3rd September, 2018

Figure 4: Flight routes of NeuroCopter at Julius Kühn-Institut. The bee on the copter knows the yellow-marked feeder on the field. The first part of the route is the same direction the bee flew from its hive towards the feeder during training phase. Flight height of the copter is 19 metres and speed is 20 km/h. Background texture of the field was created with the copter, see [29], and is overlaid on Google Earth. Images (a) and (b) show the fields and flight route of experiments of two days. The flight trajectory of the copter is depicted in red. In the lower left corner, manual start and landing are plotted, too.

4 Spike Rate

Recorded neural activity during flights with NeuroCopter is analysed. At each iteration of the experiment the copter takes nearly the same path because the route is programmed and flown autonomously. That is why the visual perception of the honeybee on the copter is nearly the same at each repetition. Thus, if this perception is correlated with the measured neuron, the neural signal of repeated routes must be similar as well. In the following, this similarity is shown.

4.1 Description

This thesis evaluates experiments from 2nd August, 2018 and 3rd September, 2018 at Julius Kühn-Institut as described in Subsection 3.3. More data was not available at the time of writing.

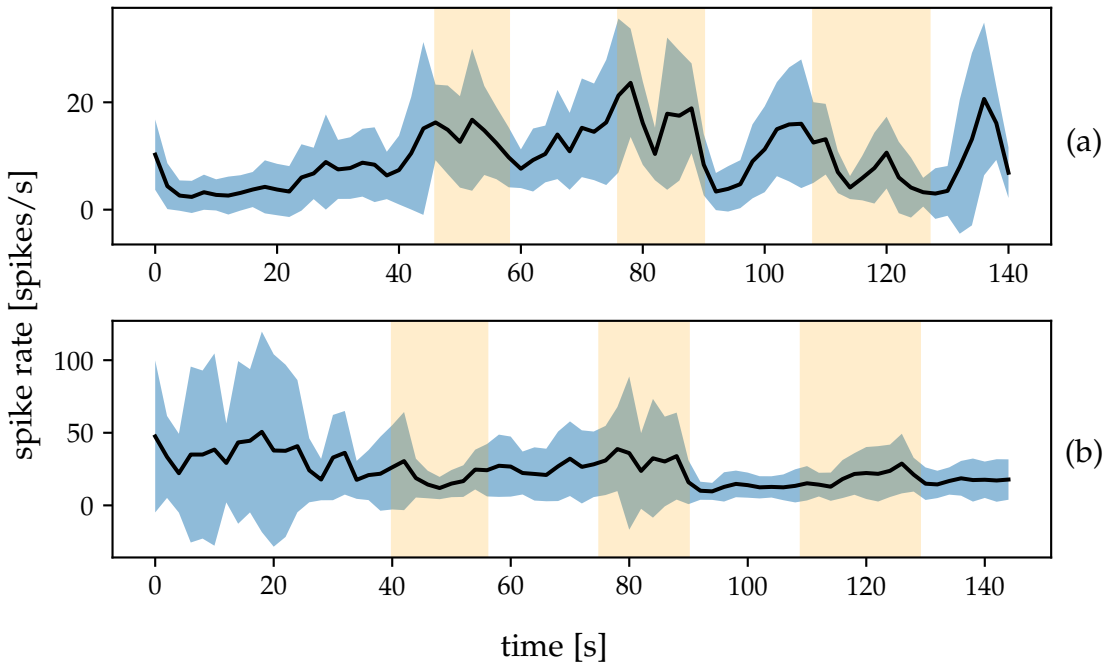

Figure 5: Mean spike rate in two second intervals of all flown rounds on 2nd August, 2018 in (a) and on 3rd September, 2018 in (b). The mean spike rate is shown in black. Standard deviation is highlighted in blue. Spike rates during turns are marked in orange.

Experiments of both days must be examined separately because the programmed route was slightly adjusted: curves flown in August, see Figure 4a, were much sharper than in September as depicted in Figure 4b. Besides, the fields were planted differently such that the bees did not see the same environment.

Each day, one honeybee was used during the field experiments. In the following, a flown path from starting to landing of the copter is considered as a flight and the route shown in Figure 4 as one round. Multiple flights were executed per day to change the battery of the copter between flights. One flight consists of three to five rounds. The waiting times varied between each round in which the copter rests at the start position of the route in the air.

The neural signal of measured neurons in the bee's brain was processed by *Spike2*, see Subsection 3.2. This software identifies and sorts spikes in the signal based on a threshold and shapes used for template matching [21]. The spike rate describes the amount of spikes per time window. In this thesis, the spike rate of the measured neurons is used in intervals of 0.1 seconds. However, this section examines the spike rate in intervals of two seconds to compare its general structures. The following sections focus on inspecting finer details. Figure 5 shows the mean spike rate of all flown rounds together with the standard deviation separated for both days. All graphs in this section are smoothed by B-splines with a smoothing condition of one.

4.2 Sliding window correlation

If there are locations on the field which the honeybee uses to navigate, the measured brain activity in consecutive rounds should be similar. The similarity of two spike rate series X and Y is measured with the Pearson correlation coefficient r. It is calculated by dividing the covariance cov of both series by their standard deviation σ [35, pp. 92 - 97]:

$$r_{XY} = \frac{cov(X,Y)}{\sigma_X \cdot \sigma_Y} = \frac{\frac{1}{n}\Sigma_{v=1}^{n}(x_v - \bar{x})(y_v - \bar{y})}{\sqrt{\frac{1}{n}\Sigma_{v=1}^{n}(x_v - \bar{x})^2}\sqrt{\frac{1}{n}\Sigma_{v=1}^{n}(y_v - \bar{y})^2}}$$

A correlation result of 0 means no linear correlation while 1 reflects a positive and -1 a negative linear correlation between X and Y. Z-normalizing is not necessary before calculating r since the Pearson coefficient is invariant to linear transformation. The Pearson correlation is computed with an implementation of the *SciPy*[4] ecosystem as well as a modified version by the author for parallel computation on the graphics processing unit with *PyTorch*.

[4] https://www.scipy.org/

If brain activities of consecutive flight rounds are similar, the Pearson coefficient of the corresponding spike rates must be greater than zero. However, the results must be compared with correlations of other parts of the flight. Otherwise, it cannot be ensured that not all parts of a flight's spike rate show these similarities. That is why a sliding window correlation (SWC) is applied [37]. First, spike rates corresponding to a flight round are selected. Next, this window is correlated with the first spike rate values of the same flight such that both series are equally long. Then, the window is shifted by one value to the following spike rate of the flight and they are correlated. This is repeated for the whole flight. Hence, the resulting correlation is one at the index which refers to the beginning of the windowed round. Moreover, results of correlations with other rounds are at their corresponding position in the correlation time series as well.

Figure 6 depicts SWCs of each round for all flights in both days in August and September. Waiting times between rounds as well as starting and landing times are differently long in each flight. They are removed before calculating SWCs to make them visually comparable between flights. Furthermore, the sum of the spike rates of each flight is used in intervals of two seconds to reduce noise. In flights a, b, and d all correlations of one round with another round are higher than with series which contain parts of the end of a round and the start of the following round. In flight e only the fourth round and in flight f the third round correlate strongly with the second round of their flights. Lastly, in flight c no round correlates strongly with another round.

The shapes of the SWCs in flights a, b and d are striking. All SWC graphs show high correlations when two rounds are correlated which is depicted at positions of green bars in Figure 6. In addition, also correlations of a round with series which start in a round and end in the next one are similar regarding all SWC graphs of a flight. These correlations are shown between the green bars. Besides, all SWCs of all flights oscillate, see Figure 6. Smooth transitions between neighbouring correlations are expected because the sliding window is shifted by one index of the flight's spike rate. Hence, consecutive correlated series with the window differ only in moved values to their predecessor index and one new successor value at the end of the series such that their correlations are similar. This results in an oscillation around highly correlated positions instead of only one peak. However, the number of oscillations between consecutive correlations of whole rounds is the same for all SWCs and all transitions between rounds in one flight. Though, this number varies between different flights.

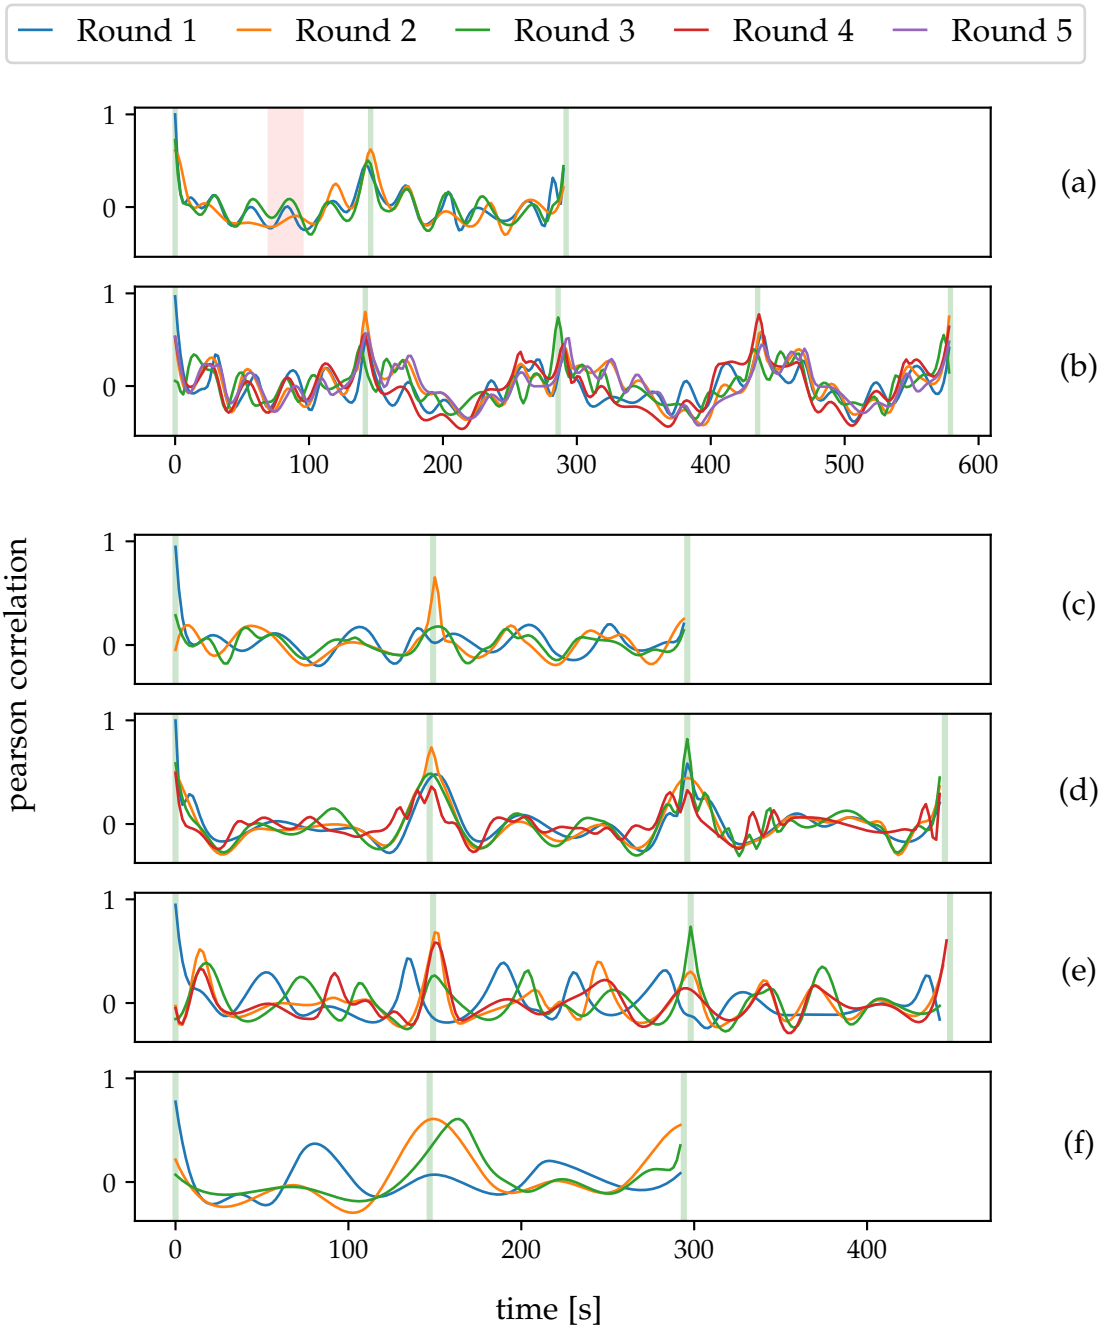

Figure 6: Sliding window correlations of all rounds of each flight. Waiting times between rounds as well as starting and landing are removed. Rows (a) and (b) show flights on 2nd August, 2018 and rows (c) to (f) show flights on 3rd September, 2018. Green bars mark starts of a flight round. All graphs oscillate between starts of rounds. One oscillation is highlighted in red in the sliding window correlation of the first round of flight (a).

In order to examine the position of the NeuroCopter at peaks of the oscillations in the SWCs a synchronised overview with a map and autocorrelation graph was developed. A slider is used to select a time of the flight which shows the location of the drone at that time on the map and highlights the corresponding position in the graph, see Figure 33.

The peaks of oscillations between rounds occur when the windowed round spike rate series is correlated with a series which starts near a turn. In August, this is often before the copter reaches the feeder the first time during a round and right after all three right turns. Peaks in September occur mostly after the first and before the third right turn. Generally, sharper turns in August lead to four oscillations in the correlations between rounds while smoother turns in September result in two oscillations in flights *c* and *d*. Flight *e* shows three oscillations between rounds that are not synchronous over all autocorrelations of that flight. The last flight *f* has only one oscillation between rounds and their peaks vary in regard to all SWCs of this flight.

4.3 Correction of different lengths of rounds

Although NeuroCopter flew the route autonomously, the total time to finish a round differs. This is observable in Figure 6 because green markings are positioned at different times between flights of the same day. They highlight start times of rounds. Hence, for example markings of the second round in flights *d* and *e* would be at the exact same time if rounds were equally long in time. These different round lengths could be caused by heavier winds which the copter cannot compensate and let it drift slightly away such that it takes longer to complete the route. The SWC graphs are calculated by correlating one fixed round with all possible time series of the same length in the flight. If the fixed round is shorter than another round, not all spike rates of the round are used for correlation. On the other hand, a longer fixed round correlates with spike rate values which include parts of the previous or next round. Even small differences in round lengths can influence the correlation results since the spike rate is used in summed parts of two seconds. Shifted oscillation of different SWCs at the same time as visible in flights *c* and *e* might be an effect of these different round lengths.

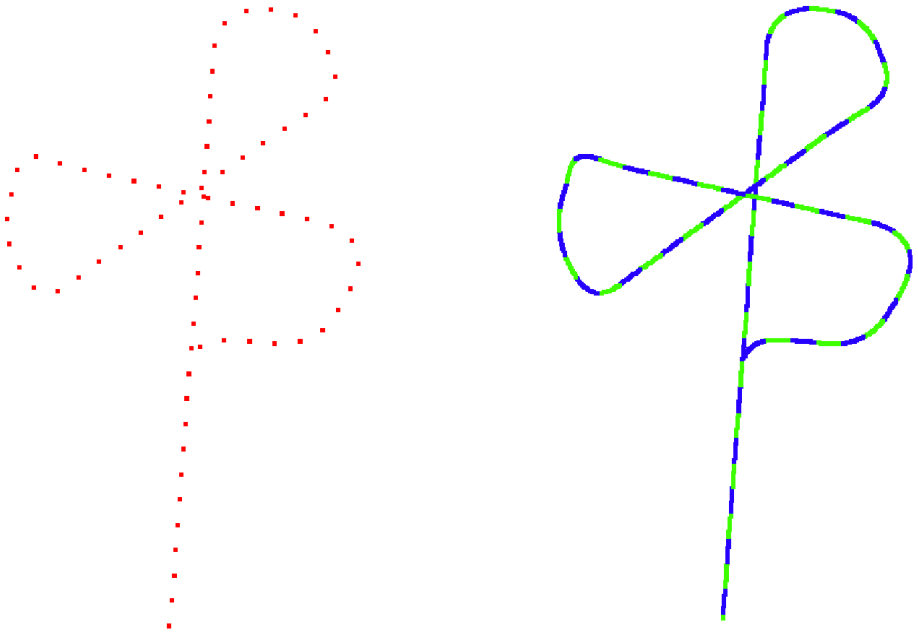

(a) Mean positions

(b) Assigned spike rate values to nearest mean positions

Figure 7: Intermediate results of compensation of different round time lengths. Image (a) shows one point every ten metres along the flight route of one round. In image (b) spike rate values of a different round are assigned to the nearest located mean points. Spike rate values are drawn at the position where they were recorded. Positions are coloured alternating between green and blue for every other mean point they are assigned to. The image serves as a sanity check to show that values are assigned properly to their nearest mean point along the route, especially in the feeder's region. The last straight return flight is not drawn to remove overlaps with the start line.

In order to improve this flaw, the spike rate was calculated dependent on flown distance of the copter rather than time. It is not possible to correlate spike rates of two rounds with different lengths because both time series must have the same amount of values to compute the Pearson coefficient. First, a distance x in which the mean spike rate will be calculated is set in metres. Next, all flight positions of the copter and corresponding spike rate values are separated after x metres. In each section a mean position is calculated at exactly half of the section length between both surrounding copter positions. An example of these mean points is shown in Figure 7a. Indices of start and end positions of a round are determined by comparing the flight position of the drone to a fixed set location. This improves the accuracy of matching rounds compared to the previous approach where start and end were determined manually.

SWCs are calculated by selecting all parts of x metres which represent the fixed round. First entries of the first part are removed if they do not belong to the round. The same is done for the last entries of the last part. Afterwards, the distance of all copter positions of another round is compared to the mean positions of the fixed round in the order of the flight. Starting with the first two mean positions, spike rate values are assigned to the first part as long as the corresponding distance to the first mean point is smaller than to the following one. Otherwise, the value is assigned to the next part and distance comparisons are continued with the second and third mean points and so forth. Figure 7b depicts the assignment of spike rate values of one round to another one.

Subsequently, the mean of spike rates in each section is calculated. This approach ensures that exactly all the spike rate values of a round are used. Moreover, the spike rates of two rounds are grouped in such a way that they are synchronised in regard to the location that triggered them.

SWCs with distance correction of all rounds for each flight are depicted in Figure 8. To compare the results with the previous approach, the distance was set to 11.5 metres which corresponds to a two second window as used in Subsection 4.2. A SWC is computed by selecting a fixed round and matching all other rounds to the fixed round with the described approach. Then, all matched rounds are joined to a time series and correlated with the sliding window of the fixed round.

Correlations of different rounds are more synchronised in flight *f*. Flights a to *d* are slightly more synchronised. However, no changes are noticeable in flight *e*.

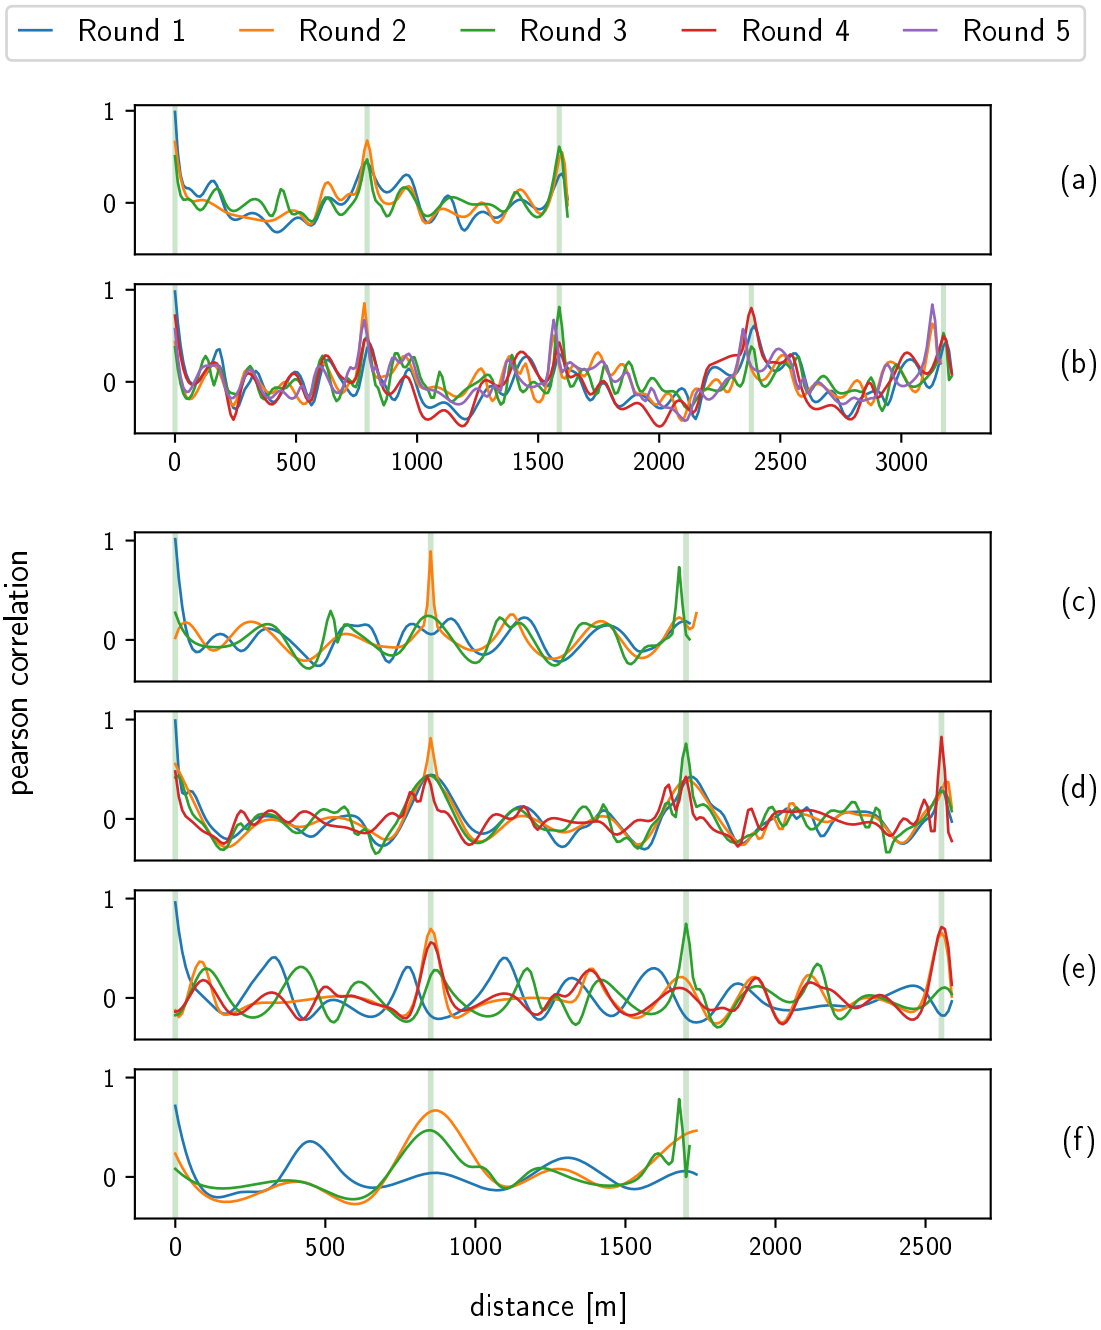

Figure 8: Sliding window correlations of all rounds of each flight with length correction. Analogous to Figure 6.

4.4 Dynamic Time Warping

The spike rates were synchronised in regards to the locations at which they were triggered between rounds in Subsection 4.3. The assumption was that as soon as specific parts of the environment are in the viewing area of the honeybee, it recognises them which results in higher brain activity. However, this might not happen at the exact same location of the bee in different rounds. For example, the bee could focus its attention a bit earlier or later at that location due to its large viewing area, see Figure 21. Hence, it is possible that the spike rate recorded in one round is a distorted variant of a spike rate of another round. This can be checked with Dynamic Time Warping which was originally used for speech recognition [34]. The algorithm stretches and compresses parts of a time series to fit another time series as best as possible.

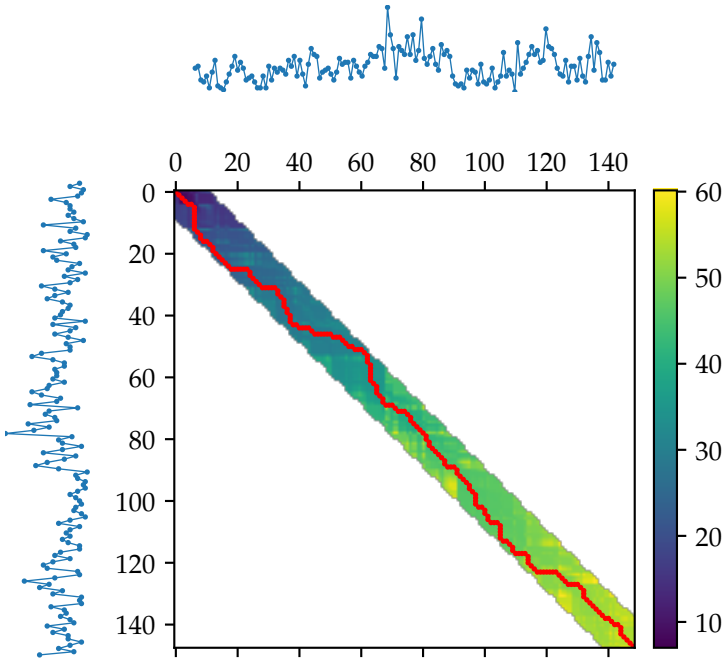

Figure 9: Calculation of shortest warping path. Indices of each axis represent indices of the corresponding time series. Dynamic Time Warping calculates euclidean distances between each point in one series and all points of the other one. Starting from the first index, cumulative distances are registered in a matrix till the end index of both series. Backtracking from the last entry provides the shortest path to the start. The depicted time series are the bee's spike rates of the first and second round in flight two on 3rd September, 2018. Not all entries of the matrix are filled because the maximum warping distance is constrained with a Sakoe-Chiba band of size ten.

Dynamic Time Warping is computed with the library *dtaidistance*[5] which was developed by the *Declarative Languages and Artificial Intelligence Research Group*[6]. The algorithm computes the euclidean distance of possible ways to warp a time series. The cumulative distances between points of both series are written in a matrix along all paths in which at least one time series is continued at each next entry. Furthermore, paths must start at the beginning and end at the last entry of the time series. In order to prevent too much warping, the maximal stretching or compressing is constrained with the Sakoe-Chiba band [33]. An example of such a matrix of two rounds with a constraint is shown in Figure 9. Lastly, the warping path with the shortest distance is obtained by backtracking. When both series are the same, the shortest path is a diagonal in the matrix. Divergences of the path from the diagonal compress or stretch a series.

In Figure 10, spike rates of the same rounds as in Figure 9 are shown with their optimal warping connections given the same constraint. Some points are matched to multiple points of the other series which stretches the signal. Compressing happens when parts are matched to the same point in the second series.

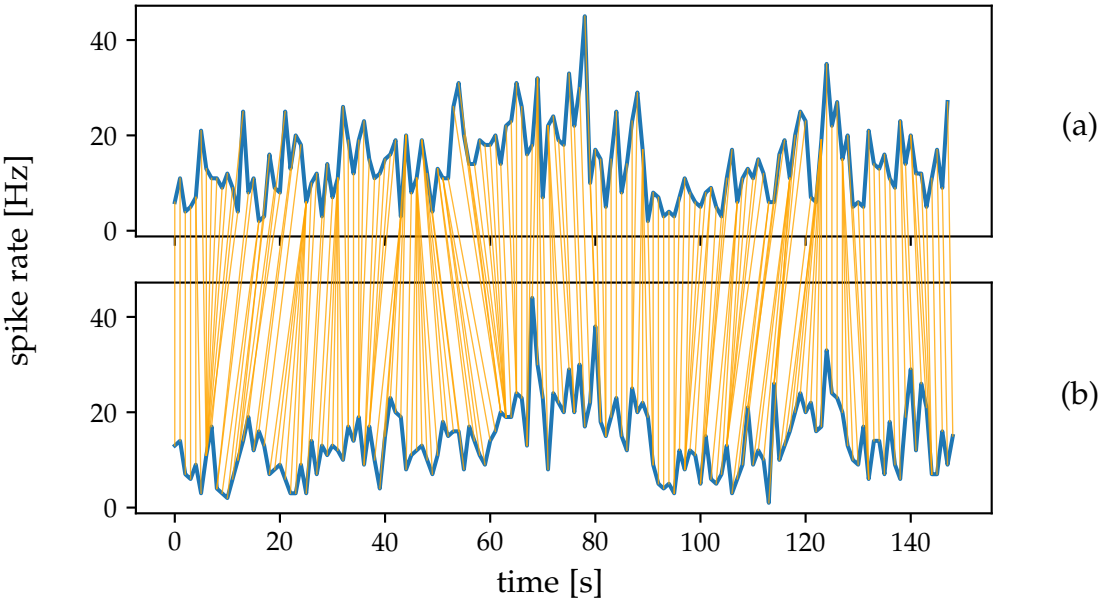

Figure 10: Visualisation of matched entries of two time series according to the shortest warping path calculated with Dynamic Time warping and constrained by a Sakoe-Chiba band. The same series and size of the Sakoe-Chiba band are used as in Figure 9.

[5] https://github.com/wannesm/dtaidistance
[6] https://dtai.cs.kuleuven.be/

Applying Dynamic Time Warping to fit spike rates of one round to another one extends or shortens the series when both rounds do not have the same length. The sliding window correlations as in Subsection 4.2 and Subsection 4.3 are modified such that spike rates of each round are warped to the round used for SWC. That is why the length of the resulting SWC graph depends on the used round whereby correlations of the same round start at different indices. Hence, these graphs are not comparable any more. In order to prevent differences, first the spike rate of each round of the flight is synchronised with the spike rate of the round used for SWC via the locations of the copter as described in Subsection 4.3. The synchronisation distance is half a metre which corresponds to mostly one spike rate value per part since the copter flies with a speed of about 5.5 metres per second and the spike rate is used in an interval of 0.1 seconds. The small distance is chosen to prevent too much adjustment before warping.

After location synchronisation all spike rates of rounds of a flight are warped to the spike rate of the round used for SWC. A Sakoe-Chiba band is used to constraint the maximal stretching or compressing to one second. Again, the sum is calculated in parts of two seconds to have comparable results with SWCs of Subsection 4.2 and Subsection 4.3. The mean values of the Dynamic Time Warping are attached in the order of the corresponding rounds. Lastly, the sliding window correlation is computed.

Figure 11 depicts SWCs with time warped spike rates. Compared to Figure 6 and Figure 8 correlations of different rounds are higher. Especially flights *c* and *e* show peaks at round starts at all SWCs. Besides, these peaks are narrower which shows that the Dynamic Time Warping does not cause higher correlations when the spike rate of one round is correlated with the warped time series of another round and parts of its previous or following round. In addition, oscillations between rounds are not as clearly visible as in previous SWCs and they are still not synchronous in flight *e*.

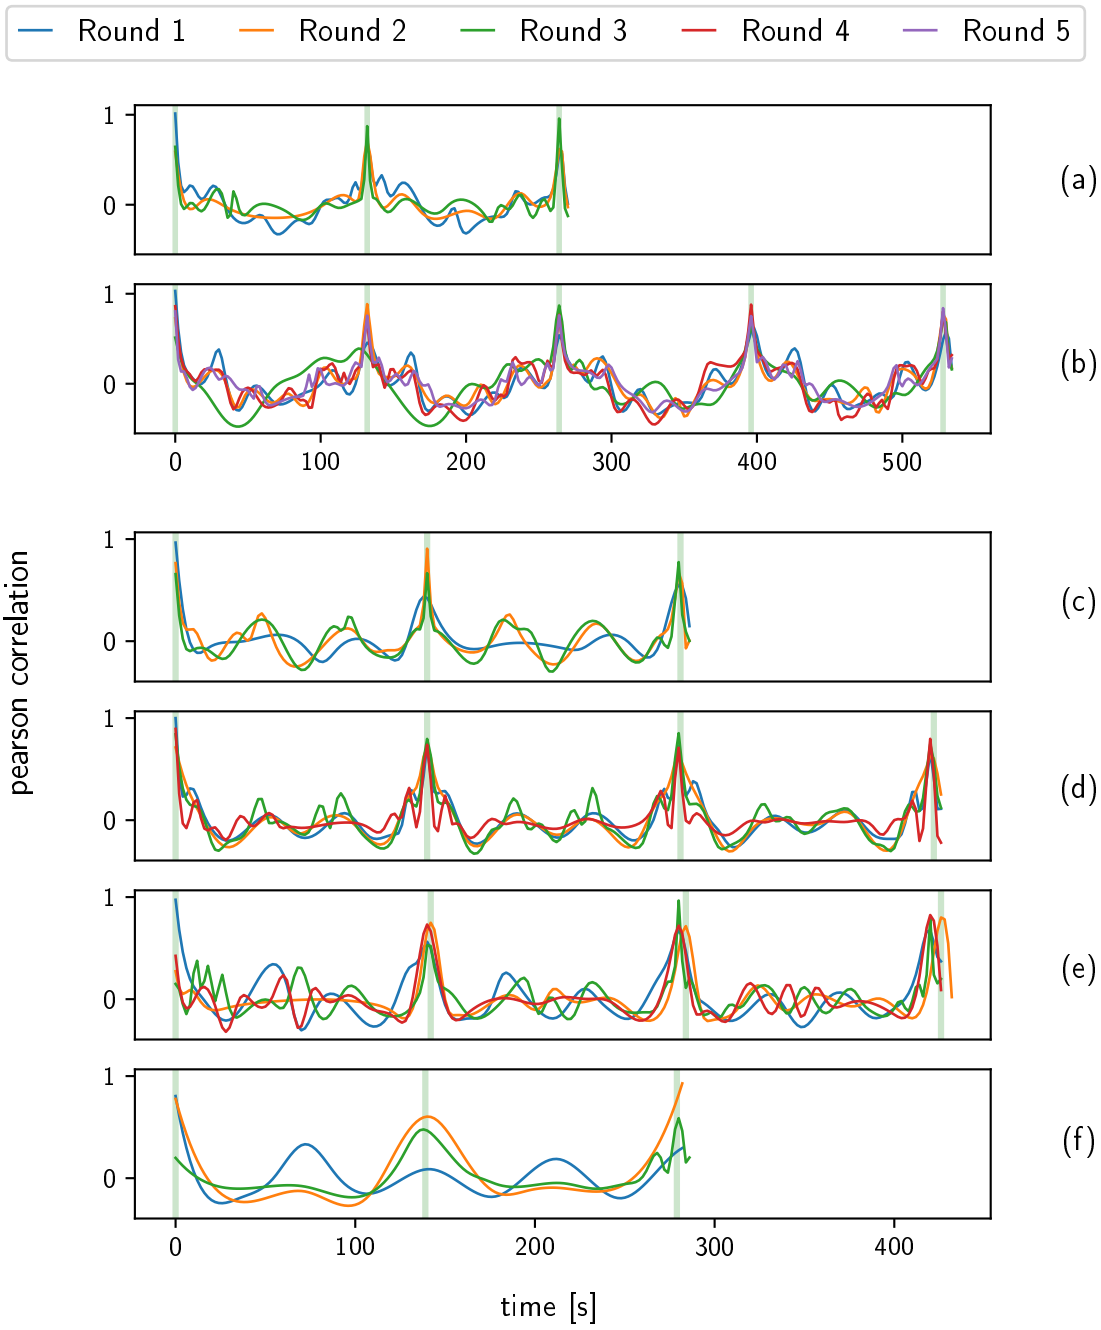

Figure 11: Sliding window correlations of all rounds of each flight with length correction and Dynamic Time Warping. Analogous to Figure 6 and Figure 8.

5 Autoencoder

In Section 4 it was shown that the spike rate recorded during flights with NeuroCopter is similar when the same path is flown successively with the same honeybee. Next, the brain activity is analysed in regard to the visual environment. The flight is simulated in a virtual three dimensional (3D) model of the experiment's area. Images are rendered from the perspective of a bee and encoded with an autoencoder. The intermediate outputs unite parts of the image. Correlating these outputs with brain activity should indicate which part of the bee's view area triggers neural spikes. In addition, high correlations could suggest that the neural spike rate can be determined by the autoencoder.

5.1 Description

In order to check which visual structures lead to high brain activity in a honeybee while flying, every part of the animal's visual perception needs to be examined. For this task, a convolutional neural network is suitable since it learns features of images, from low- to high-level patterns, and is based on the structure of an animal's visual nervous system [15]. However, supervised learning by training a neural network to determine the spike rate according to an image was not feasible because there was not enough flight data with NeuroCopter available. That is why unsupervised training is needed. Hence, a β-variational autoencoder [17] with convolutional layers is used to encode the image. Subsequently, the encoded parts can be correlated with the spike rate.

The autoencoder is implemented with the deep learning platform *PyTorch*[7]. The architecture of the used autoencoder is depicted in Figure 12. First, the dimension of the input image is reduced by convolutional layers to a latent representation. This part is called encoder. Then the input is recreated by upsampling and convolution which is referred to as decoder.

5.2 Generating bee images

Training the autoencoder requires a lot of images. On top of that, the training data should not be included in the images which will be used for actual analysis to avoid that the model only memorises known data instead of generalizing during training. That is why it is not sufficient to take images of videos from flights which are recorded by the camera on NeuroCopter during experiments.

[7] https://pytorch.org/

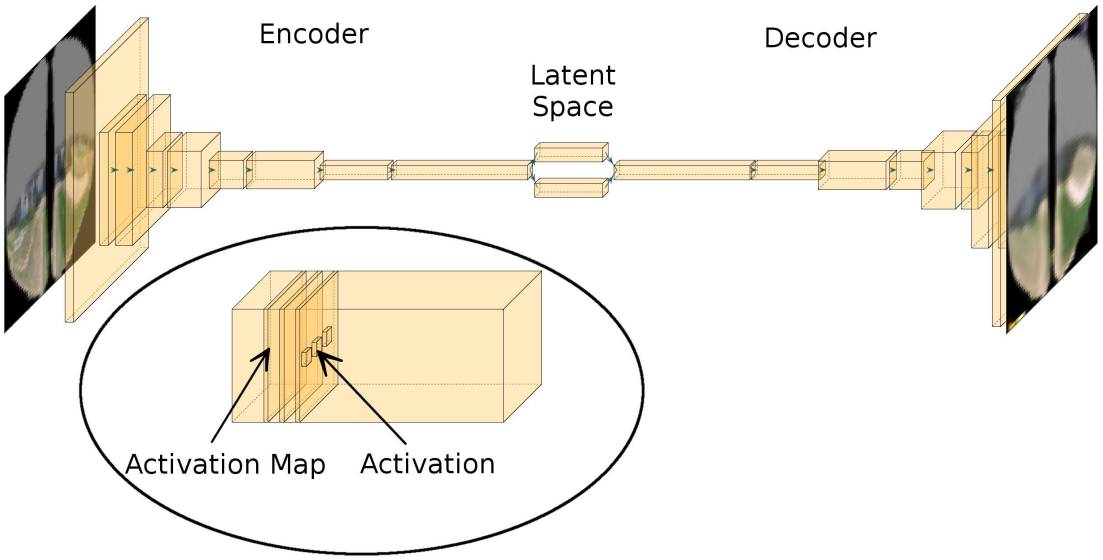

Figure 12: Architecture of the autoencoder model which is used to gather visual information and correlate it later with brain activity. The autoencoder consists of an encoder, a latent space and a decoder. This graphic shows the size of each resulting layer after a weight-normed convolution followed by a Scaled Exponential Linear Unit (SELU) activation function. In addition, the decoder uses upsampling to reach the input size. Each layer has multiple activation maps with activations as shown in the highlighted layer. From encoder to latent space as well as from latent space to decoder fully connected layers are used. The latent space represents a distribution of the input images from which the decoder reconstructs the input images. The notions *activation map* and *activation* are marked in an enlarged layer. (Created with *PlotNeuralNet*[8])

On the other hand, additional flights with the copter could be done to capture images for training. But this approach has many drawbacks for later analysis: the recording of the camera on the copter was not synchronised with the recordings of the brain activity. Hence, manual selection of start points in both recordings is error-prone. Moreover, the frame rate of the camera would restrict the possible intervals of spike rate and flight images. On top of that, high temperatures during experiments led to artefacts on the recorded videos which affect the analysis.

That is why a virtual 3D environment was used instead of camera footage of NeuroCopter. This makes it possible to create a lot of random images for training. The flight of the copter can be simulated by recorded GPS data. This approach provides the opportunity to adjust the frame rate of flight images such that different spike rates can be used.

[8] https://github.com/HarisIqbal88/PlotNeuralNet

In addition, the 3D model enables generating images in perspective of bees with the library *bee view* from Johannes Polster [31]. These bee view images could not be created by transforming video frames of the copter because a much wider field of view is required as will be shown in Figure 22.

The copter is able to take multiple overlapping photos of the field automatically and save their GPS position and camera orientation such that a 3D model of it can be constructed with *Agisoft PhotoScan* as described in [29]. This aerial mapping ensures that the 3D model reflects the actual field at the time of the experiment as close as possible. However, the 3D model of the field is not sufficient because honeybees also see the surrounding area during flights with NeuroCopter. Hence, another 3D model was created of the environment. Since aerial mapping is too complex for this task, Julian Petrasch generated an environmental model with *ERDAS IMAGINE*[9] by using NASA Shuttle Radar Topographic Mission[10] (SRTM) which provides elevation data and satellite images from the Landsat program[11]. In order to keep the size of the whole model as small as possible Julian Petrasch generated a 3D model of the near environment surrounding the field with high resolution and a model which covers large parts of Berlin with lower resolution.

Figure 13: Virtual 3D model which consists of the fields at Julius Kühn-Institut (1) where experiments with NeuroCopter took place as well as the near (2) and far (3) surrounding environment. The fields have a size of about 50.000 m^2 and the whole model of about 220 km^2.

[9] https://www.hexagongeospatial.com/products/power-portfolio/erdas-imagine
[10] http://srtm.csi.cgiar.org/
[11] https://landsat.gsfc.nasa.gov/

The author of this thesis combined the models of the field and environments with *Blender* [12] such that the transitions between these models look natural. A part of the combined model is shown in Figure 13 with the field of the experiments in the centre. The 3D model does not contain a texture for the sky because no panoramic photos were taken on the same day as experiments were conducted. Photos of another day were not used since different weather conditions as well as other positions of the sun and clouds might influence the analysis. Overall, the aim was to find locations on the field. That is why the sky of bee view images is grey-coloured as shown in the input image of the autoencoder's architecture in Figure 12.

5.3 Training

The autoencoder was trained with renderings of the 3D model shown in Figure 13. Using the *bee view* library [31] these images represent views from a perspective of bees. The training set consists of 100.000 bee images which were taken at random locations on the field. The height was set randomly of a normal distribution with at least three metres. Moreover, the direction vector of the virtual camera was chosen such that the yaw angle is taken randomly from a uniform distribution and the pitch angle is drawn from a normal distribution with mean zero and standard deviation 0.05. This assures that all directions are covered and that the training images are taken from angles in which the bee could been tilted while flying with NeuroCopter.

The autoencoder was trained for 100 epochs. Figure 32 in the appendix shows the loss during training and validation. An example of input and output image is depicted in Figure 12.

5.4 Correlation with brain activity

The trained encoder of the autoencoder extracts features of bee images which are used as the visual perception of the bee. That is, the decoder of the autoencoder is not needed any more. In order to check if the brain activity represents these visual features during flights, bee view images are generated in the 3D model along the flown path and with the copter's orientation during experiments. The flights can be simulated accurately with the help of recorded GPS coordinates and tilt angles of the copter. Next, these images are processed by the encoder in the correct order of the flight. Each activation of every activation map is tracked, see Figure 12, such that there is one time series with the measured values of all bee images during the flight for each activation.

[12] https://www.blender.org/

Finally, the time series are correlated with the spike rate of recorded brain activity. High correlation values indicate a relation between visual perception and neural activity in the brain of the bee. The Pearson correlation coefficient is used as described in Subsection 4.2 to measure the linear correlation between activation and spike rate time series. The spike rates in intervals of 0.1 seconds of the first and second flights of 3rd September, 2018 are analysed with the autoencoder in this section.

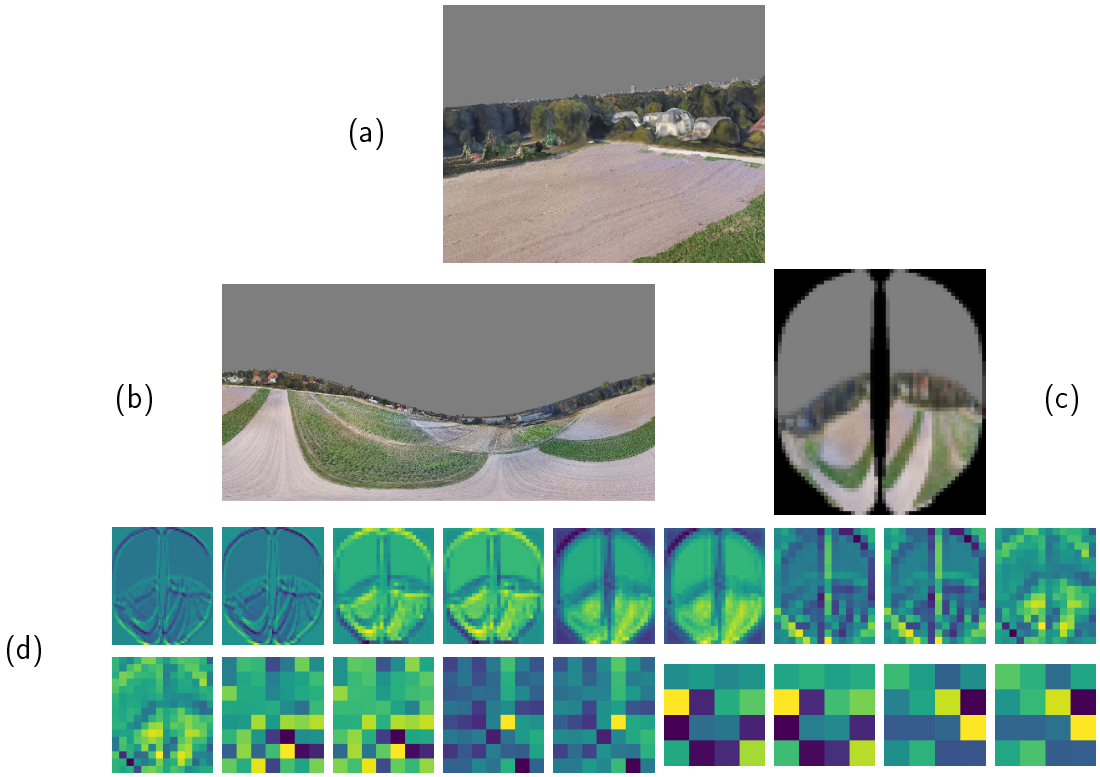

Figure 14: Example of the encoding of a bee view image with an autoencoder. Image (a) shows the field from the position of a bee during the simulated flight of NeuroCopter in the virtual 3D environment rendered with a pinhole camera. A panorama rendering from the same position is depicted in image (b) which shows the large field of view of a bee. The animal perceives the environment at the same position as shown in image (c). This image is encoded with the autoencoder. One activation map per layer is depicted in (b). Every other activation map shows the output of weight-normalised convolution and a SELU activation function.

5.4.1 Round analysis

First, it is examined if there is one activation in the encoder part of the autoencoder which changes its value similar to the bee's spike rate during a flight round. Image *a* of Figure 14 shows the field from the position of a bee during the flight. Images b and c show the field of view of the bee and its perspective at the same position as in image *a*. The encoder processes the bee's perspective which produces multiple activation maps. One of them is depicted in *d* of Figure 14 for each layer. Activation maps of first layers contain the contour of the field as visible in the bee image while following maps are increasingly blurry. Outputs of last layers do not show shapes of the input image.

Each pixel of these activation maps corresponds to one activation. Bee views were rendered with a frame rate of ten images per second such that there is one spike rate value for each image. Each image is encoded by the autoencoder in the order of the flight. For all frames all activations are tracked. That is, there are as many recorded activation values per entry of all activation maps as bee view images are processed. Figure 15 shows the correlations of activations with the spike rate in one round. The sky is not correlated with the brain activity since its colour is always grey. It does not have any texture, for example clouds, as visible in image *a* of Figure 14. Maximum and minimum Pearson correlations of positive and negative 0.4 are located at the horizon. The contours of the fields show negative correlation values.

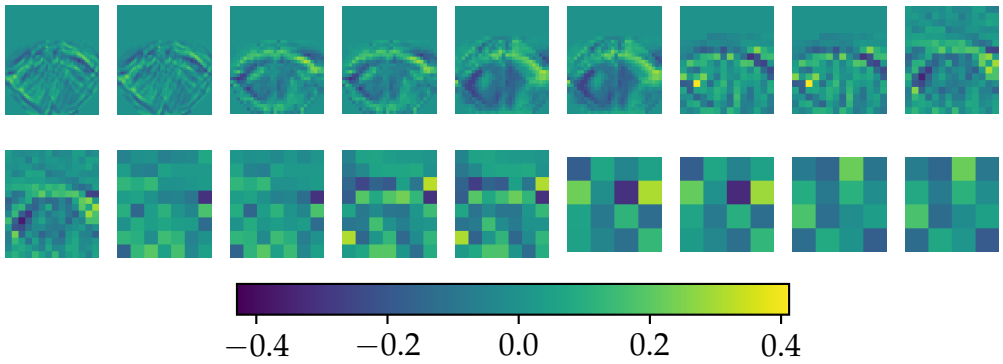

Figure 15: Correlation of activations with spike rate during one flight round. Changes of all activation values are tracked individually while the encoder processes bee view images along the flight route. The resulting activation series are correlated with the corresponding spike rates of the round. One activation map of each layer of the encoder is shown with the correlation values of each activation with the spike rate.

Figure 15 shows correlations of only one flight round. Likewise, correlations of activations with the spike rate are calculated in the remaining rounds of the first and second flights on 3rd September, 2018. The distribution of all resulting correlations separated by layers of used activations for all rounds of the first flight on 3rd September, 2018 is depicted in Figure 16. The histogram of the second flight in Figure 34 in the appendix shows similar results.

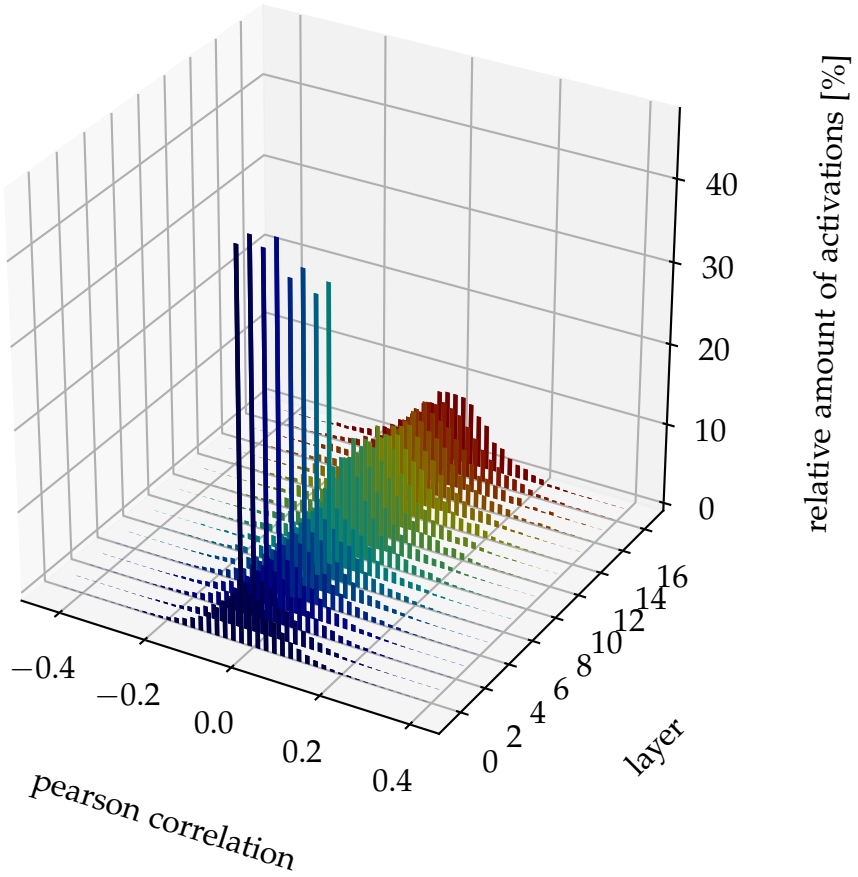

Figure 16: Distribution of all correlations of activations with spike rate during all rounds of the first flight on 3rd September, 2018. Correlations are separated by layers of the corresponding activation. Histograms of each layer are normalised to show the relative distribution of correlation values per layer.

The correlation values are normally distributed with a mean of zero. In layers one to eight between 28% and 41% of all activations do not correlate with the spike rate because their Pearson value is zero. These are activations which represent the static sky as well as the corners of the image which are always black due to the distorted bee view projection. As visible in Figure 14 the transitions between image corners, sky and field become fuzzier in each layer. This affects the correlations in Figure 15. In the first layer the corners and the sky areas have assigned correlations with zeros while the areas decrease till layer nine. The subsequent layers do not contain these areas which have the same value during all frames.

Across all layers, the maximum Pearson correlation is reached at about 0.4 in the first flight and 0.3 in the second flight. Hence, it is not visible that activations with high correlations are located in a specific layer.

Figure 15 and Figure 16 show that there are activations which partly correlate with the spike rate. Next, areas of the bee view images are visualised which are used by activations that have high correlations. The goal is to find parts of the bee's vision that trigger similar activation values to the spike rate. While Figure 15 gives an overview of locations of activations with high correlation values in activation maps, it does not show which areas of the input images are mostly used for these activations.

For a given activation a the area of the input image of the autoencoder can be computed in the following way. The indices of row r and column c of a in its activation map are stored. The variable *size* describes the amount of surrounding area around the entry at (r, c) and is initialized with zero. Then, for each layer l starting from the previous layer of a to the first one:

Increment *size* if l is convolutional layer with stride 1
: The encoder has nine convolutional layers as visible in Figure 12. All of them use a kernel size of three and padding of one. Every other convolutional layer starting from the first one moves the kernel with a stride of one. The resulting activation maps have the same size. That means, each entry of the produced activation map is computed by convolving the entry at the same row and column as well as the eight surrounding pixels with the kernel. That is why r and c are not changed and *size* is increased by one.

Double r, c, *size* and increment *size* if l is convolutional layer with stride 2
: The remaining convolutional layers use the same kernel and padding sizes but the kernel is moved with a stride of two. These convolutions halve the size of the resulting activation maps. Hence, after convolution each entry is computed with entries around the activation at row $2 \cdot r$ and column $2 \cdot c$. If *size* is already greater or equal to one the entries

surrounding the activation at row r and column c must be considered as well. Their row and column indices are also doubled. That is why the area of the input image used for convolution, which is described by *size*, is multiplied by two. In addition, convolution with kernel size three again increases *size* by one.

Do nothing if l is a SELU activation function

All convolutional layers are followed by a SELU activation function. These layers are not depicted in Figure 12 but are shown in image *b* of Figure 14. Since activation functions do not use surrounding entries of activation maps for calculation, none of the variables r, c and *size* is changed.

Lastly, all pixels of the input image which are used during convolutions to compute a can be marked. One pixel is located at row r and column c while the surrounding pixels are within the range of *size* in positive and negative direction row- and column-wise.

Results of this approach are shown in Figure 17 for each round in both flights. All rounds are shown with the areas used by activations with the 20 and 5000 highest correlation values to the spike rate. Regarding the 20 highest correlations, the corresponding activations are mostly calculated with pixels on the left and right part of the horizon as visible in round one and two of the first flight and three and four of the second flight. The first two rounds of the second flight use pixels in the sky and upper-field area. Round three of the first flight is the only round in which the left area of the field and the right-bottom part are used. Drawing areas used for the 5000 highest correlations reveals an intersection at the horizon in all rounds except the third round of flight one.

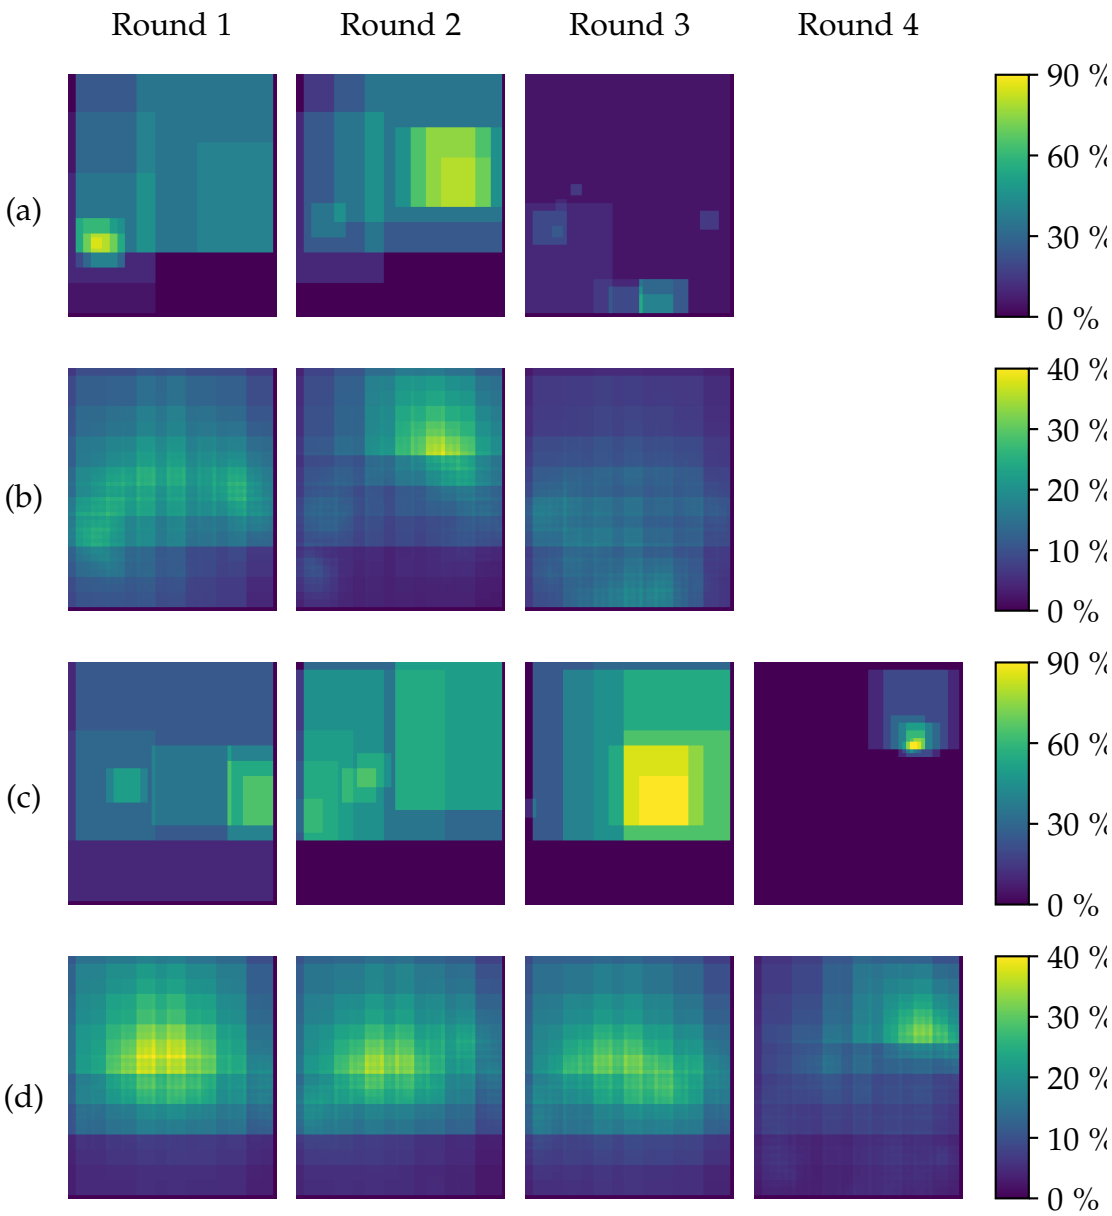

Figure 17: Used regions of input images by activations with highest correlations per flight round to the spike rate. Each image has the same size as the input bee view images of the encoder. The colour of each pixel denotes the amount of activations which are calculated, among others, with this pixel. The amount is relative to the overall number of depicted activation regions. Row (a) shows used input regions of activations with the 20 highest correlations in the first flight on 3rd September, 2018. Activations with the 5000 highest correlations are used in row (b). Rows (c) and (d) show results with analogous settings for the second flight on 3rd September, 2018.

5.4.2 Local analysis

Instead of correlating activations with the spike rate of a whole round, correlations are computed in two second sections. Hence, for every activation there are multiple correlation values. This enables finding activations which might conform to the spike rate only in several parts of the flight route. Nevertheless, activation with more sections with high correlations are preferred. That is why the sum of all correlations per section for each activation is calculated to compare them. Figure 18 depicts the pixels used by activations with the 100 highest correlation sums. Again, the horizon in the bee's vision is important area in most rounds. In addition, parts of the field are marked in all rounds. The sky is used in the first two rounds of the first flight and the second and fourth round of the second flight.

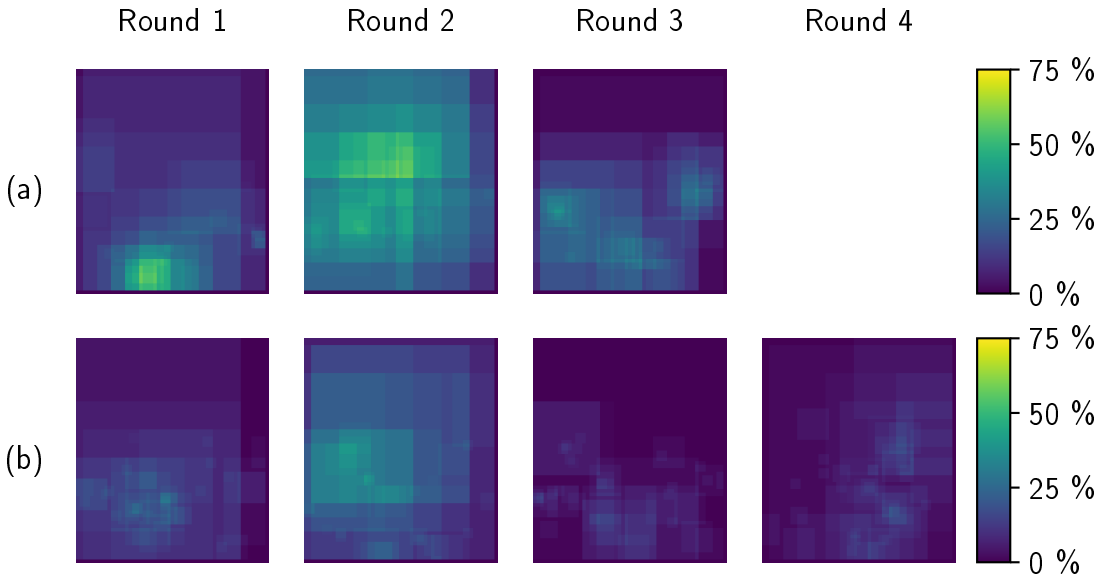

Figure 18: Used regions of input images from activations with the highest sums of correlations in two second sequences per flight round. For each round the activations with the 100 highest correlations sums are used. Row (a) shows results of the first flight and row (b) for the second flight on 3rd September, 2018.

5.4.3 Detection of best activations

Previous sections focus on locating activations which match the spike rate best in one round. However, the goal is to detect activations which are changed similar to the spike rate in all rounds. Analogously to Subsection 4.2, a modified sliding window correlation (SWC) is used to measure similarity of rounds for an activation and the spike rate. The whole flight is used for SWCs including waiting times between each round as well as the start and end. For a given activation sequence a of a flight the SWC with a sliding window of the first round is computed as follows:

Correlate the first round of a and the spike rate in sequences of two seconds
As in subsubsection 5.4.2 correlations of an activation and the spike rate are split in sequences of two seconds. Correlations which belong to the first round form the *base series* used for SWC.

Calculate sliding window correlations of the *base series* with all correlations of two second sequences
For each index i a correlation series, named *indexed series*, is computed in the same way as *base series* such that the *indexed series* has the same amount of correlations as the *base series*. The index i starts from zero and ends at the last possible index of a which has sufficient successors to create a correlation series that has the same length as *base series*. The *base series* is correlated with each *indexed series*.

This approach compares correlations of one round in two second sequences with correlation sequences of other rounds. Moreover, correlations of with overlapping parts of waiting times and flight rounds are computed. In the following, correlations of the first round are always used as the sliding window. The result of this modified SWC is shown in Figure 19 for the activation which has the highest sum of correlation values in two second sequences in the first round. That is, it has the highest correlation value with the local analysis approach described in subsubsection 5.4.2. The SWC equals one at the index of the beginning of the first round because the first round is correlated with itself. However, correlation sequences of the first round do not reach a higher correlation with correlation sequences of another round than with overlapping rounds and waiting times.

Ideally, the correlations of different rounds should correlate with a high Pearson value while the correlation of a round with overlapping parts of rounds should yield a low correlation result. As described before, this does not apply to the SWC shown in Figure 19 although the activation with the highest sum of all correlations in two second sequences with the spike rate in the first round is used.

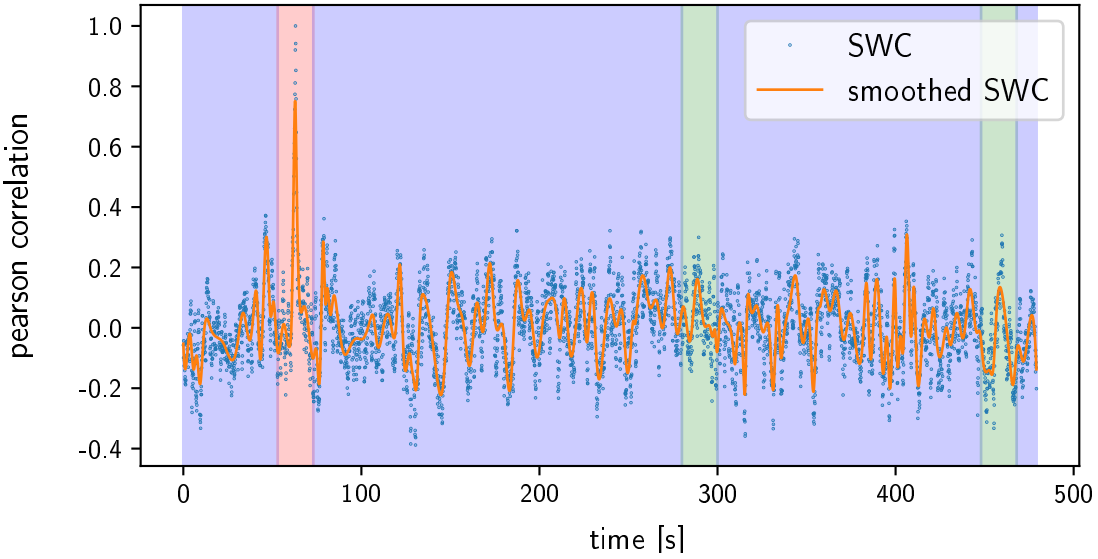

Figure 19: Sliding window correlation of correlates of the first round in two second sequences of one activation with the spike rate. One activation of the autoencoder is correlated with the spike rate in two second sequences over the whole flight. Then, the correlation series which represents the first round is used for SWC over all correlations. Blue dots represent the SWC time series and the orange graph smooths this signal. Distributions of sections highlighted in blue and green are compared with the Kolmogorov-Smirnov and Mann-Whitney rank tests. The section in red is not considered because it correlates the first round with itself.

In order to find activations which fulfil this condition best, SWCs of all activations have to be compared. While the same task could be done by visually inspecting graphs of SWCs of the spike rate for different rounds and flights in Subsection 4.2, it is not possible in this case with activations. Overall, the encoder part of the autoencoder contains about 138.000 different activations which all result in distinct SWC graphs. Hence, this task must be automated.

SWCs of all activations were computed by correlating the correlations of the first round with all other possible correlation sequences as described above. All SWCs are divided into two sets. One set contains correlations at the index of beginnings of rounds after the first round and correlations up to ten seconds before and after these rounds start. Correlations which belong to this set are highlighted with a green background in Figure 19. The correlation of the first round with itself which is highlighted in red is left out. All remaining correlations belong to the second set, they are highlighted in blue.

Next, the distributions of the sets containing blue and green correlations are compared. The Kolmogorov-Smirnov test [35, pp. 401 - 405] and the Mann-Whitney rank test [35, pp. 361 - 363] are used with the corresponding implementations of *SciPy*. They check whether both sample sets are drawn from the same distribution which results in a high p-value up to one. Since the goal is to find activations whose blue and green highlighted sets are different, the p-values should be near zero. Figure 20 shows pixels of the input image that are used for activations with the 20 and 5000 lowest p-values in both tests and for both flights. Results of both tests are similar. Activations with the lowest p-values use pixels at the horizon on the right in both flights. Besides, the right part of the horizon is visible when pixels of activations with the 5000 lowest p-values are marked in the first flight while the left part of the sky is highlighted in the second flight.

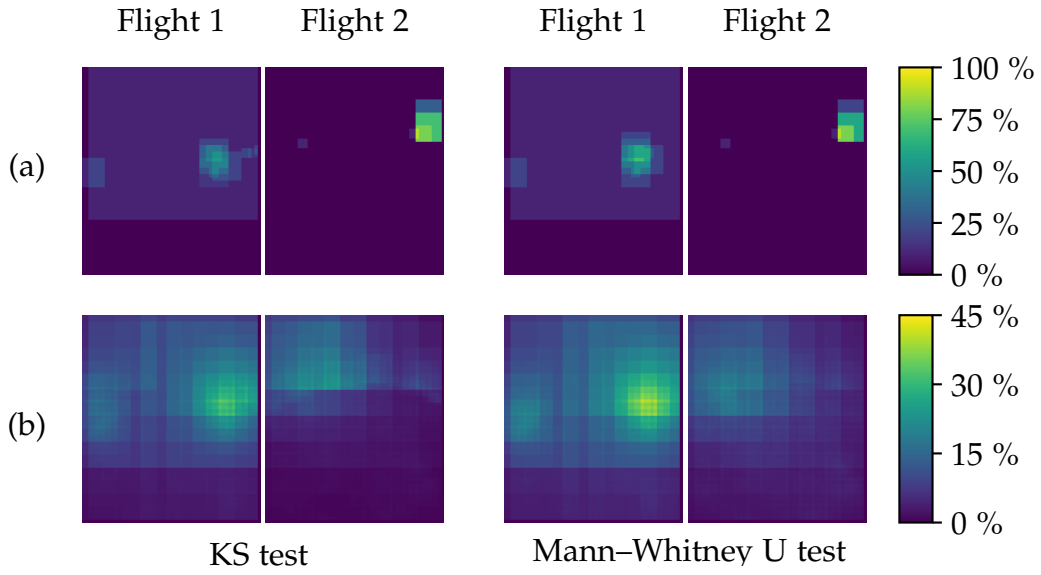

Figure 20: Used regions of input images from activations with highest differences between round correlations and correlations with shifted signals. Correlations of all activations with the spike rate are calculated in two second sequences. Correlates which represent the first round are used for sliding window correlation with all correlations of the flight. The distribution of correlations of the first round with the following ones is compared to the distribution of all other correlations while leaving out the correlation of the first round with itself. Row (a) shows input regions of activations with the 20 greatest differences and row (b) with the 5000 greatest differences of both distributions measured separately with the Kolmogorov-Smirnov and Mann-Whitney rank tests. Results for the first and second flight on 3rd September, 2018 are depicted.

6 Field Mapping

In Section 4 it was shown that the neural brain activity of repeated flight rounds is similar and Section 5 describes an approach for correlation with the visual perception of the honeybee. However, it is still not identified which visual structures trigger brain activity. That is why in the following the brain activity is plotted on the field.

6.1 Description

The 3D model as described in Subsection 5.2 is used to determine which parts a honeybee sees during flights with NeuroCopter. The library *bee view* [31] was modified to return coordinates on the field which are in viewing range of a bee at a particular point.

Figure 21: Close up photograph of a honeybee. The eyes are composed of about 5500 ommatidia which bees use to perceive an image of the environment. The enlarged extract of the bee's eye makes a part of the ommatidia noticeable. The bee's field of view is over 313° large, only the region behind their thorax is not visible. [36] (Photograph from USGS Bee Inventory and Monitoring Lab[13])

[13] https://www.flickr.com/photos/usgsbiml/34717512800/in/album-72157664305903459/

To generate bee views, the library uses ray casting for each ommatidium which is a part of the bee's eyes, see Figure 21. That is, a line is positioned in the virtual simulation at the locations of all ommatidia and aligned in the viewing direction. Then, the first intersection with another 3D model is calculated [32]. This way, coordinates on the field are determined which the bee sees. The colours of hit locations on the 3D model are then used to generate the bee image. These hit locations are used to create the mapping images in this section.

Figure 22 shows the mapping process of a bee view on the map. First, the virtual bee is positioned at a location of the flight route and aligned like the NeuroCopter at that position. The bee view is created as shown in image *a* of Figure 22. Each pixel in the bee view corresponds to the view of one ommatidium. Hence, each pixel in the bee view relates to one location in the 3D model.

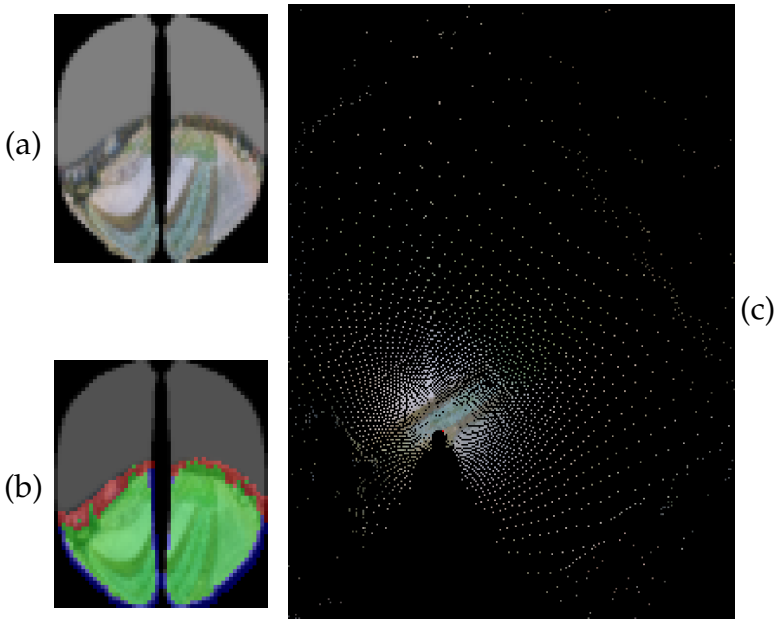

Figure 22: Mapping of a bee's vision at one position. Image (a) shows the view of the field in a perspective of a bee. In order to draw this perspective onto a map, the location of each pixel in the 3D model is calculated. Only pixels which show a part of the field are used for mapping which are marked as green in image (b). Red pixels show parts of the environment outside the field. They are not used for mapping as well as blue pixels which lead to artefacts after mapping. The locations of hit ray casts of green pixels are shown in figure (c) from a top view. Each mapped pixel is assigned the same colour as in the bee image.

Only pixels which show a location on the field are used for mapping. These pixels are highlighted in green as depicted in image b of Figure 22. The x and y coordinates of each location are marked on a new image relative to the coordinates of the upper-left corner of the field. The z coordinate which describes the height is not used as visible in image *c* of Figure 22. Thus, the new image shows the viewing area of a honeybee at one location from the top. In order to visualize the relation between the bee view image and the map, pixels on the map are assigned the same colours as their corresponding pixels in the bee view image.

When the copter with the bee moves the viewing area moves as well. Figure 23 depicts the viewed area during a flight of one line as well as one round. In these maps the shape of the field as well as different parts of it can be distinguished. Moreover, it shows that the composed viewing area of the honeybee after the first line is about as detailed as after the whole round.

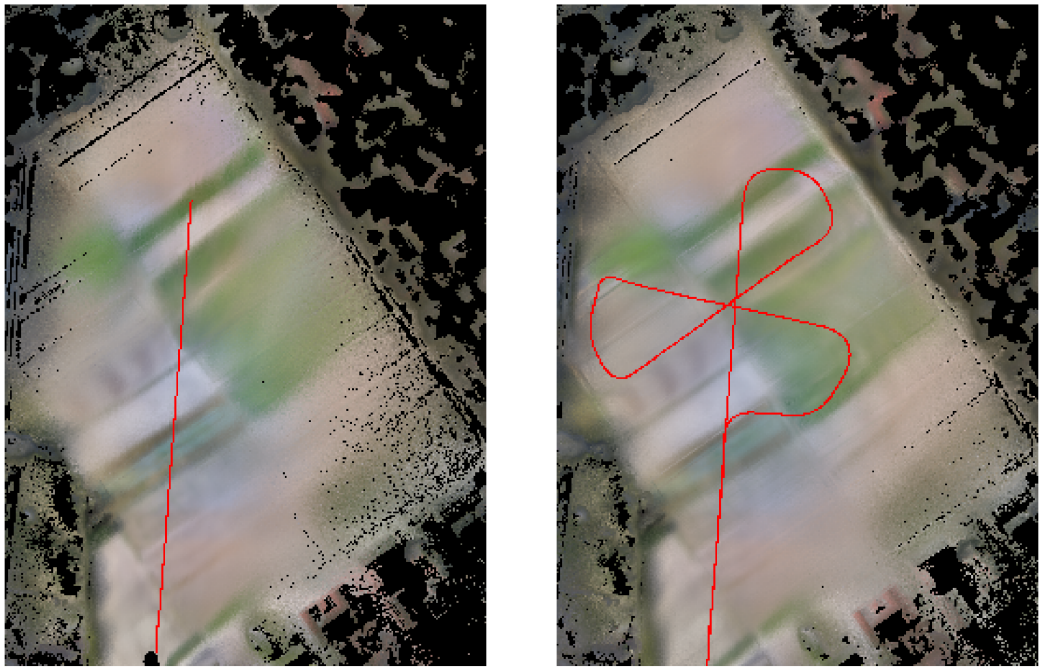

(a) Viewing area of a straight flight

(b) Viewing area of a flight round

Figure 23: Viewing area of a honeybee from top in the simulated 3D environment. All pixels in (a) are viewed during the first line on the route while (b) shows all viewed pixels after flying one round. The used flight trajectory is shown in red. The viewing area of the bee is projected on the map in 0.1 second intervals. The mean is calculated of pixels on the map with multiple assigned colours.

Compared to the field recorded by the copter in Figure 4 the mapped bee views show clear separations between field parts over which the bee flies. Other field parts are blurry such as the top of the field and the lower-right corner.

In the following subsections, different time series are plotted on the map at the hit locations of ray casts to visualise where the values arose. In Subsection 6.2 one ray cast per ommatidium is used while Subsection 6.3 shows maps with multiple ray casts per ommatidium as it is used in the *bee view* library to generate bee views.

6.2 Mappings

Multiple time series are plotted on the field. The flights are simulated as described in Subsection 5.2. During a flight some coordinates on the field are hit multiple times by ray casts. To visualise the map either the mean, median, standard deviation or variance of the values of these coordinates are calculated. Moreover, separate plots without values during flight curves are shown because the spike rate is already known to be higher than during lines [29] which might distort the result of the map.

6.2.1 Activations of autoencoder

First, the bee images processed by the autoencoder in Section 5 were mapped. Correlating single activations with the spike rate of the bee's brain activity lead to high matches with activations which use pixels of the bee image in the region of the horizon instead of the field. Pixels in this area alternate between showing the field and the background due to the different distortions of the bee image when the alignment of the bee is changed. These changes result in higher activity at pixels on the horizon than on the field. The mapping process described in this section is not influenced by rotations of the bee because for each bee image only coordinates on the field are used. This is shown by green pixels in image *b* of Figure 22. Therefore, mapping the values of activation maps of the autoencoder on the field could give new insights.

Figure 24 shows mapping of one activation map of the autoencoder's first layer after the first line and after one round. As in Figure 23, parts of the field are separated. Edges and paths between have lower activation values than pixels belonging to field parts. Again, areas over which the bee flies are sharper than parts of the field which are further away. Since paths and edges are the only information which are encoded in these maps, activation maps are not further examined.

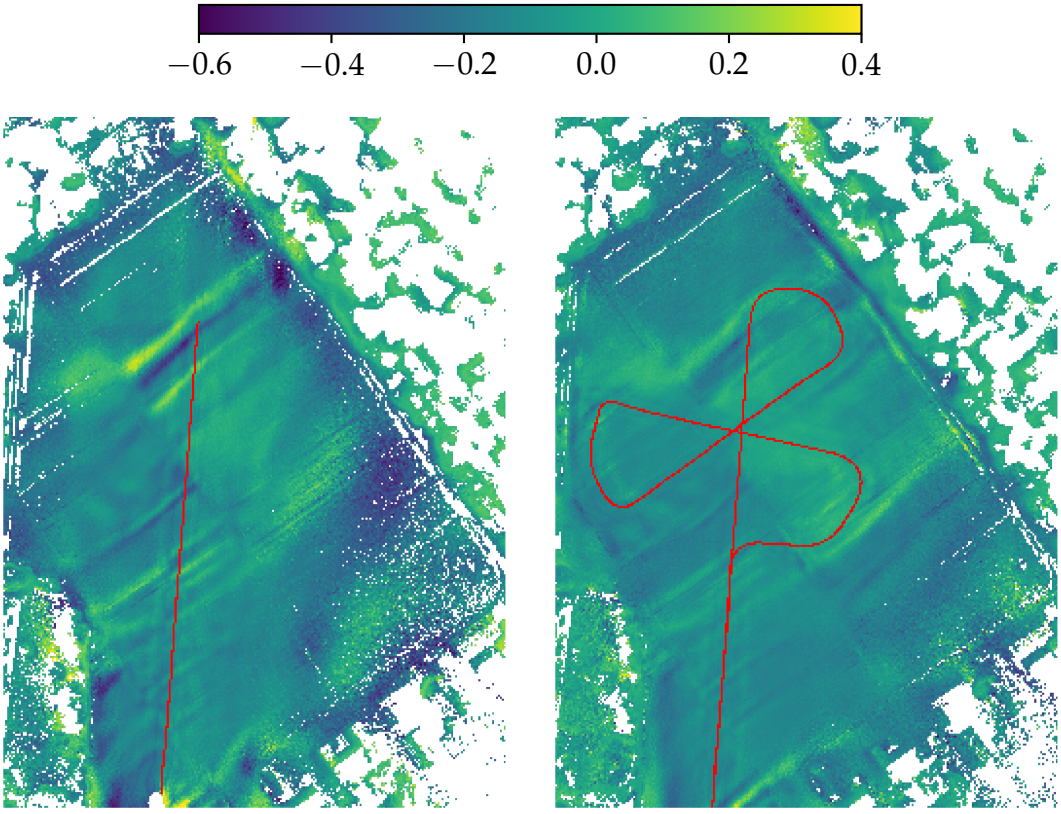

(a) Activations on map after a straight flight

(b) Activations on map after a flight round

Figure 24: Mapping of activation maps of an autoencoder. The viewing area of a bee is mapped in a 0.1 second interval. Each bee image is encoded by the autoencoder described in Section 5. A resulting activation map of the first layer is projected on a map in the viewing area of the bee. The mean is taken of pixels with multiple mapped activation values. Image (a) shows mapped activations after a line flight and (b) after one round as marked by the red flight trajectory.

6.2.2 Spike Rate

Another approach is to map the brain activity of the honeybee. The position of the bee in the simulation is updated ten times per second. The spike rate of the measured neurons is used in the same interval. At each position, all pixels in the viewing area of the bee on the map are assigned the corresponding spike rate value at that time. In comparison to drawing the spike rate on the flight trajectory, the mapping technique highlights which locations on the field trigger high spike rates instead of showing the positions of the bee with high brain activity. In addition, the maps show if the locations lead to high spike rates from different or only one direction.

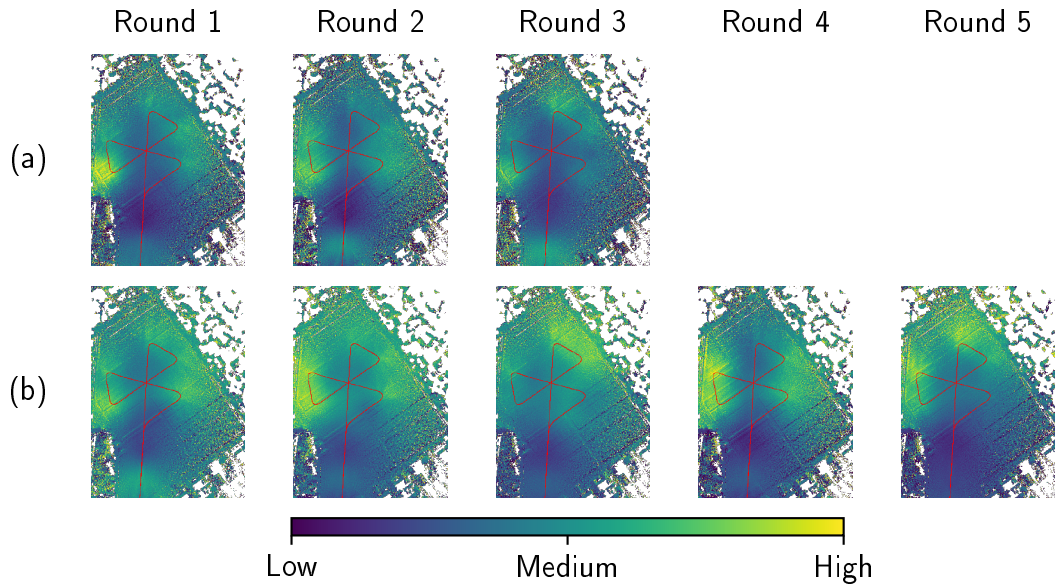

Figure 25: Mapping of the spike rate of a bee during flights on 2nd August, 2018. Each pixel on the map which is in the viewing area of the bee is assigned the spike rate at that time. The position of the virtual bee is updated ten times per second along the flight path of the recorded NeuroCopter. Accordingly, the used spike rate interval is 0.1 seconds. At the end, the mean is calculated of spike rate values assigned to each pixel. The values of all maps are normalised to compare locations which trigger high brain activity. Each row represents one flight. The flight trajectory is marked in red.

Spike rates of all rounds flown in August are mapped in Figure 25. The mean is calculated of pixels with multiple assigned spike rate values during mapping. The spike rates generally decline from the first to last round during a flight. Moreover, the spike rate altitudes vary across flights. That is why the values of each map are normalised between zero and one such that locations with high and low spike rates relative to each round can be compared between all flights.

Each round has two to three locations with high spike rates. All of them are at turns. The route has three large turns. The first and second turns have a location with high brain activity in all rounds while the third turn does not have such a location in round three of flight a as well as round three and five of flight *b*. Start and end sections of the route have low spike rates in all rounds.

The sharp turns always trigger high spike rates in comparison to the rest of the round. The flight route in September was improved by flying smooth turns. Figure 26 depicts mapped spike rates of all flights flown in September. The route adjustment can be compared by the flight trajectory drawn in red. Again, the mean is taken of multiple assigned spike rate values to one pixel on the map and all resulting values are normalised.

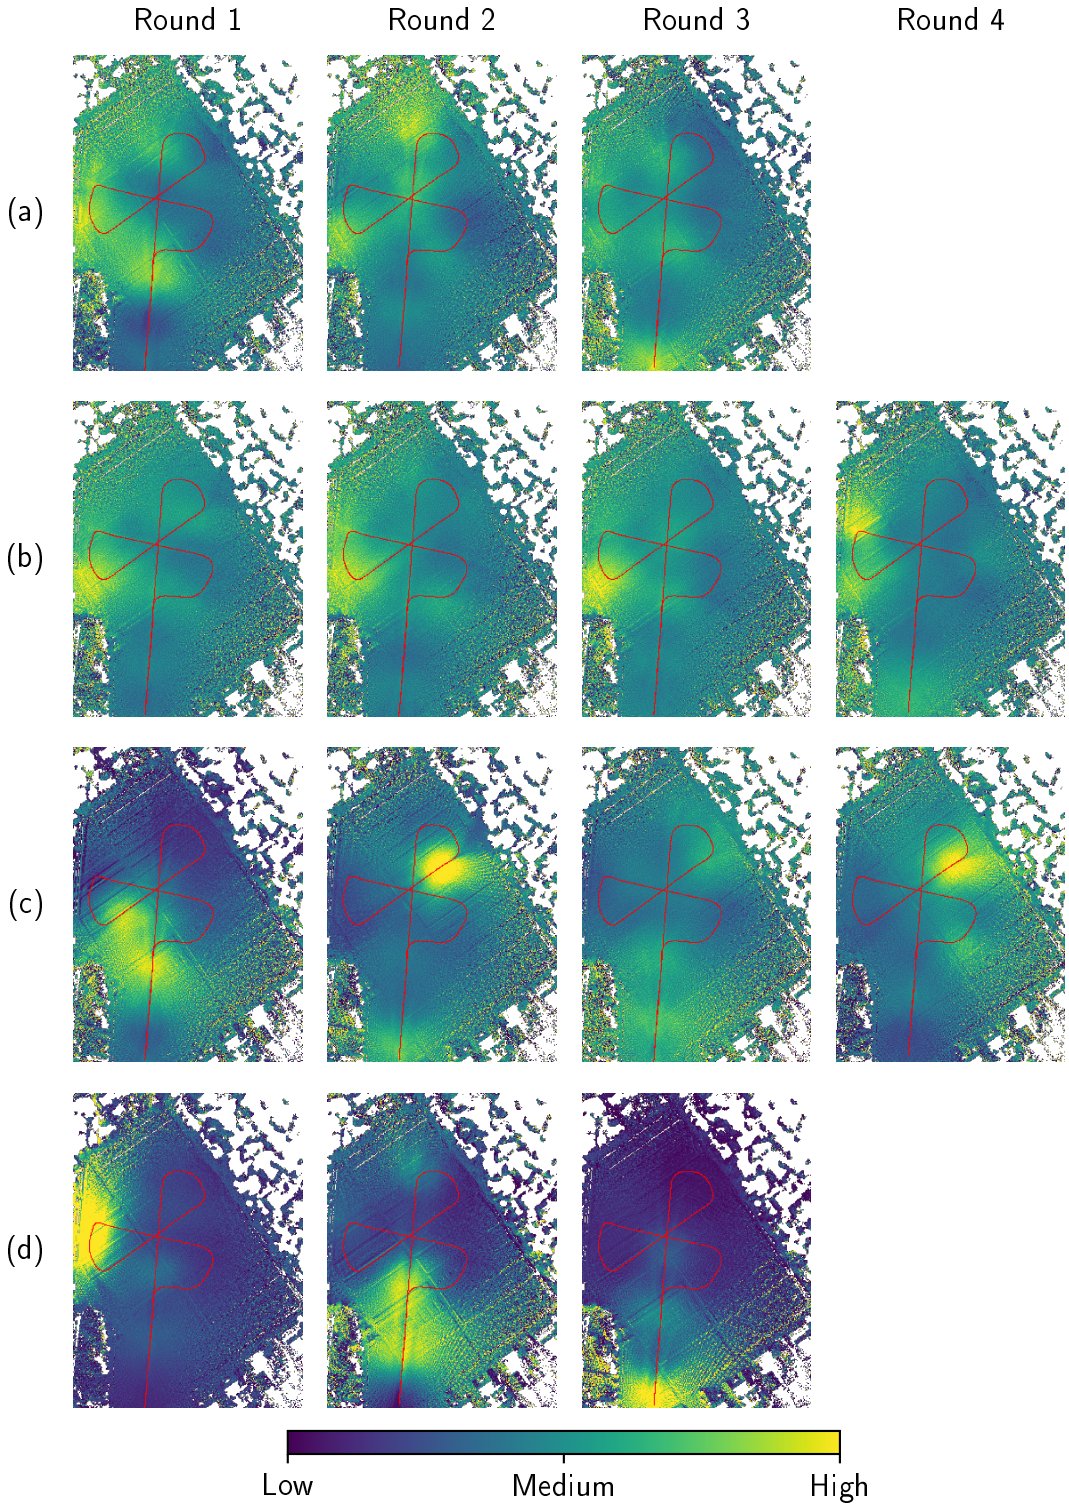

Figure 26: Mapping of the spike rate of a bee during flights on 3rd September, 2018. Continuation of the same mapping technique as in Figure 25 with flights flown in September.

In September, a location with a high spike rate appears frequently in the second turn as well. Most notably, this is the case in the first two rounds of flight *a*, all rounds of flight *b*, and the first round of flight *d*. Except of the second round of flight *a* and the second and forth rounds of flight *c*, the other two turns do not have high spike rate locations. The last both rounds of flight *d* have high spike rates in the start and end sections of the route.

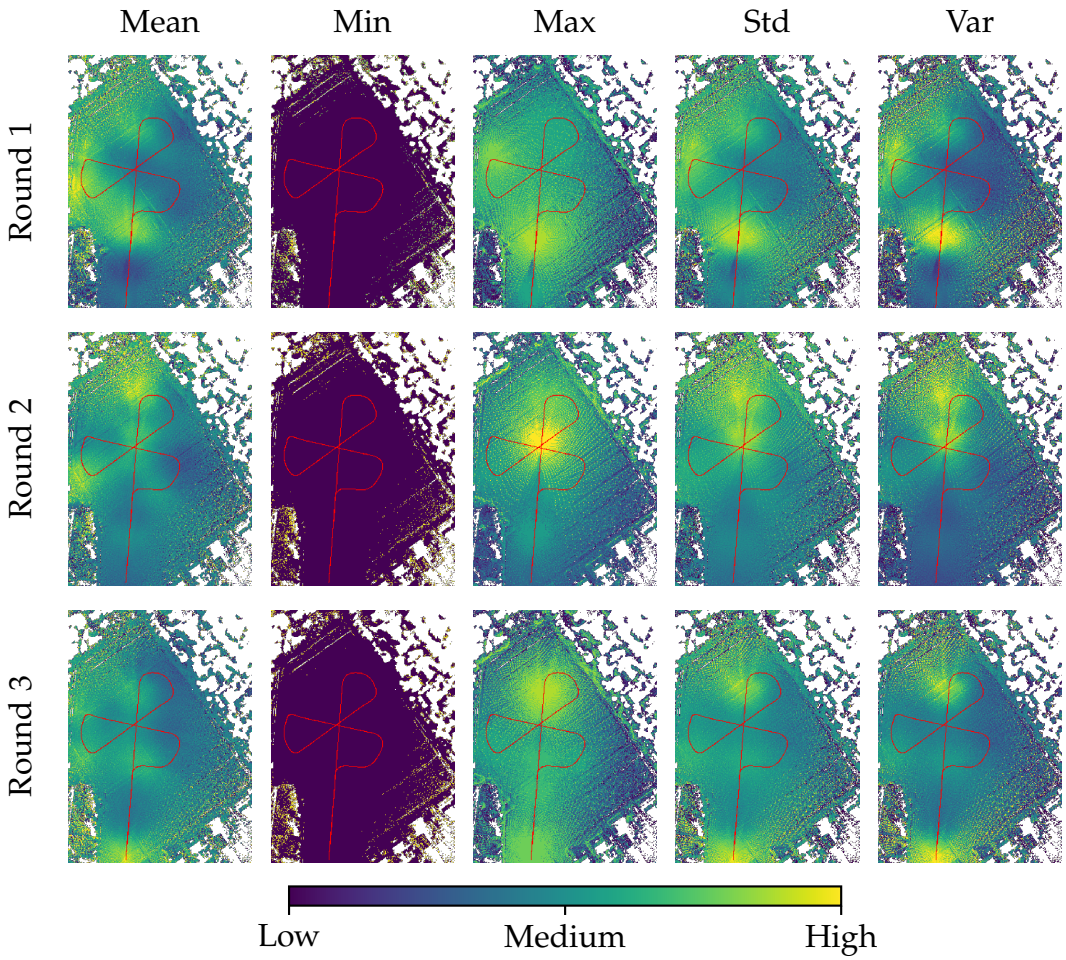

Figure 27: Application of statistics on mapped spike rates of the first flight on 3rd September, 2018. Each row shows mappings of a round where pixels on the map with multiple assigned spike rate values are calculated differently. The first column shows mean values as depicted in Figure 26. The following columns show maps with min and max values as well as standard deviation and variance.

Most maps with projected spike rates are influenced by turns such that other parts of the field have relatively small values. However, it is already known that turns trigger high spike rates. Therefore, the following approaches are executed to reveal further parts of the map which also have a high spike rate assigned in comparison to the rest of the field.

First, overlapping assignments of spike rates are depicted in other ways next to the mean in Figure 27. The first flight in September is analysed to use the route with smoother turns. Each map of the three rounds is shown with minimum and maximum values of assigned spike rates to pixels as well as with standard deviation and variance. In all three rounds, the minimum spike rate assigned to nearly each pixel is zero. Only pixels in the right corner of the field have a minimum spike rate of one. This shows that the majority of pixels on the field are assigned multiple spike rate values during the flight route since most pixels have a higher value than zero on the mean maps. The maps with maximum, standard deviation and variance have similar locations with high spike rates in the same round. In all three rounds, that location is on the first line between the start and the first turn: in the first round before reaching the feeder, in the second round at the feeder and in the third round after flying over it.

Another approach to suppress known locations with high spike rates in turns in order to reveal other location which trigger brain activity is to remove turns. Mappings of the spike rate of all flight rounds flown in September without turns are shown in Figure 28. Per round 4 lines are mapped: from start to the first turn, between the three turns and from the last turn to the end as shown by the red trajectory. Although locations with high spike rates during turns are not mapped, the remaining maps stay nearly the same in comparison to mappings with turns in Figure 26. Exceptions are round two of flight *a* which has a high spike rate area around the feeder without turns. In addition, round one of flight *d* shows such a location on the first line before reaching the feeder.

Overall, no consistent locations which trigger high spike rates besides turns could be revealed. There are rounds with high activity marked near the feeder. For example, round two in Figure 27 shows the maximum spike rate around the feeder and the first two rounds of flight *d* in Figure 28 show high brain activity before reaching the feeder. But these locations do not trigger a high spike rate in all rounds.

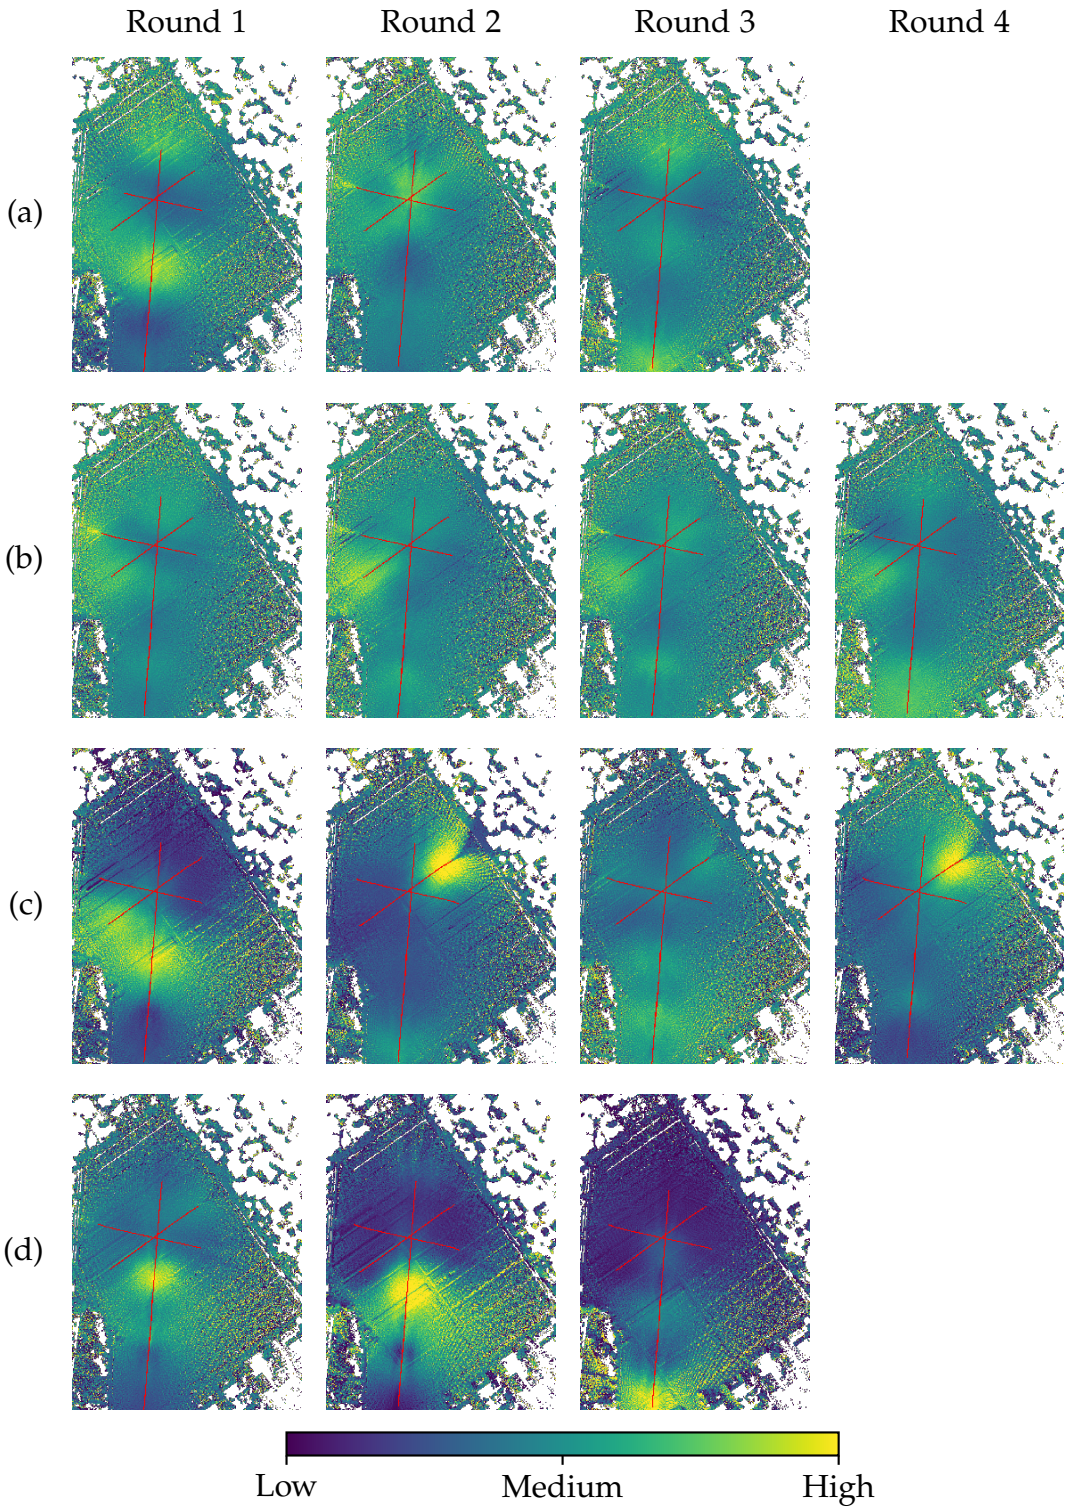

Figure 28: Mapping of the spike rate of a bee during flights without turns on 3rd September, 2018. The same approach is used as shown in Figure 26 without turns.

6.3 Multiple samples per ommatidium

The *bee view* library casts multiple rays per ommatidium and calculates one colour which is displayed at one pixel of the bee view image. In Subsection 6.2 the hit location of only one ray cast per ommatidium is used for mapping. Using the same amount of ray casts as the *bee view* library results in more detailed mappings. The viewing area with multiple ray casts in Figure 29a is sharper compared to image *c* in Figure 22. Figure 29b shows the mapping of an activation map of the autoencoder with multiple ray casts.

Calculating mappings with the higher amount of ray casts requires a higher memory consumption and is more time consuming. Moreover, there are no advantages besides sharper mappings. That is why only one ray cast was used for the mappings in Subsection 6.2.

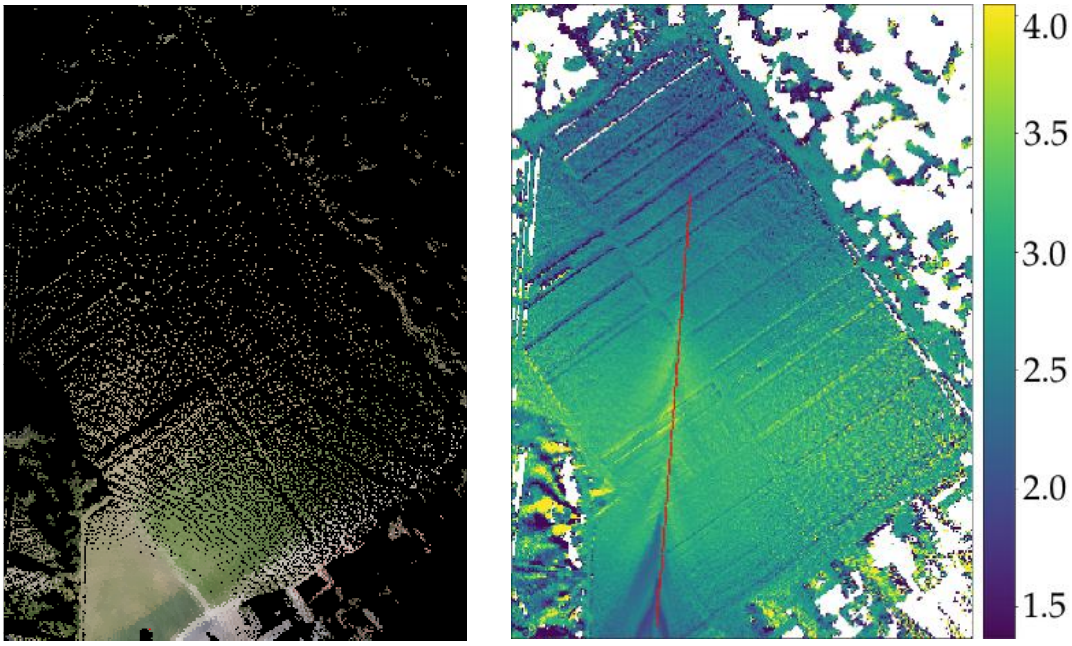

(a) Viewing area at one position

(b) Activations on map after a straight flight

Figure 29: Example images which make use of multiple ray casts per ommatidium. Picture (a) shows the viewing area of a honeybee at the first position of the route and (b) shows the mapping of activations after flying the first line of the route which is marked in red.

7 Evaluation

Previous sections analysed the spike rate of brain activity of two honeybees flying with a quadcopter. These bees knew a feeder that was placed on the field over which the copter flew as described in Subsection 3.3. The measured neuron is located in an area of the bee's brain which is responsible for processing of visual information, see Figure 2. That is why it was examined if the spike rate reflects changes in the visual field of the animal. In particular it was assumed that flying to and over the known feeder triggers high brain activity. Besides, field structures and landmarks might have an influence on the spike rate because honeybees use them for orientation, see Section 1. In the following, previous results of this thesis are evaluated in regard to these assumptions.

First it was shown in Section 4 that spike rates of successive flight rounds with the same animal are similar. That is the basis for further analysis because the copter flew autonomously along a programmed route such that the bee had roughly the same visual perception in consecutive rounds. Therefore, spike rates must be similar as well if the measured brain activity of the bee is related to its vision. Sliding window correlations (SWC) in Figure 6 of all rounds of each flight show that most spike rates of different rounds are correlated. However, some rounds do not correlate with another and some SWCs of different rounds in the same round are shifted. As a potential issue it was identified that some rounds differ in flight lengths which could be caused by compensating wind. But synchronizing spike rates of every round to the locations where they were triggered in Figure 8 resolved only shifting of the SWCs. Additionally applying Dynamic Time Warping with a Sakoe-Chiba band constraint of two seconds which allows warping up to one second in future and one in past in Figure 11 improved correlations between previously uncorrelated rounds. Hence, spike rates of different rounds of the same flight are not equal but similar.

Some SWCs of the same flight have similar correlations even between round starts in form of oscillations. Peaks of these oscillations mark correlations of a round with a time series starting at a turn of a round and ending in the following round. Since there are no peaks when the copter flies straight, turns are prominent in the spike rate. High spike rates in turns were already discovered in [29]. These findings suggest that the spike rate reflects the motor speed and acceleration of NeuroCopter instead of actual brain activity. Motor speed as well as acceleration always change when turns are flown compared to straight lines. In fact, neural signals in the honeybee's brain are very small and interfered by noise of the copter's motors. That is why shielding and grounding is used on the copter [29]. Figure 35 shows that in some rounds the spike rate and motor speed is correlated. Though, this is not the case for all rounds.

There are also rounds in which the spike rate is not correlated with motor speed at all. The same applies to the results of correlations of acceleration with spike rate in Figure 36. On the other hand, changes in motor speed and acceleration impact the visual perception of the connected honeybee. Since the measured neuron processes visual signals, partly correlations are not deniable. Consequently, measured spike rates do not only reflect motor speed or acceleration.

After it is known that the bee's visual perception and its neural activity are each similar across multiple flight rounds, two approaches for correlating vision and brain activity were conducted. The idea of the first approach is to process the visual perception of the honeybee with an autoencoder in a virtual 3D model of the environment where experiments with NeuroCopter were executed. The encoder part of the autoencoder received bee view images along the flown route of the copter. Activations of the encoder were correlated with the spike rate. Correlations were computed for whole rounds in Figure 17 and sequences of rounds in Figure 18. Moreover, SWCs of these correlations were calculated to find activations with similar correlations in consecutive rounds which are different from remaining correlations in Figure 20.

Pixels of the input image are visualized which are used for convolutions to calculate best fitting activations in each scenario. Highlighted areas of the input images are in all cases around the horizon. Often, the right part of the horizon is marked. Input areas of activations with the 5000 highest correlations are shaped as the field in a bee's perspective during a straight flight as shown in Figure 30a. In many rounds activations with the 20 highest correlations use pixels at the right part of the horizon. This can be explained by the distorted bee view image during turns as shown in Figure 30b. The flight route contains three big right turns and one small left turn at the end. While the whole horizon jitters due to small tilts of the copter, the right part of the horizon is affected the most. Pixels of the input image in this area change often between showing a part of the field and background. This results in high activity in activation that use these pixels. That is why the autoencoder approach detects flight turns instead of locations on the field which trigger spike rates such as the known feeder or other landmarks such as paths on the field.

A drawback of the simulated 3D environment is its RGB colour space. All pixels are represented of red, green and blue ratios. However, bees have ultraviolet-sensitive, blue-sensitive and green-sensitive photoreceptors [13]. To capture the field in the appropriate colour space for bees, the copter would need a camera which is able to take images with ultraviolet light. That is why bee view images of the 3D simulated environment do not completely imitate the visual perception of a bee in the real world.

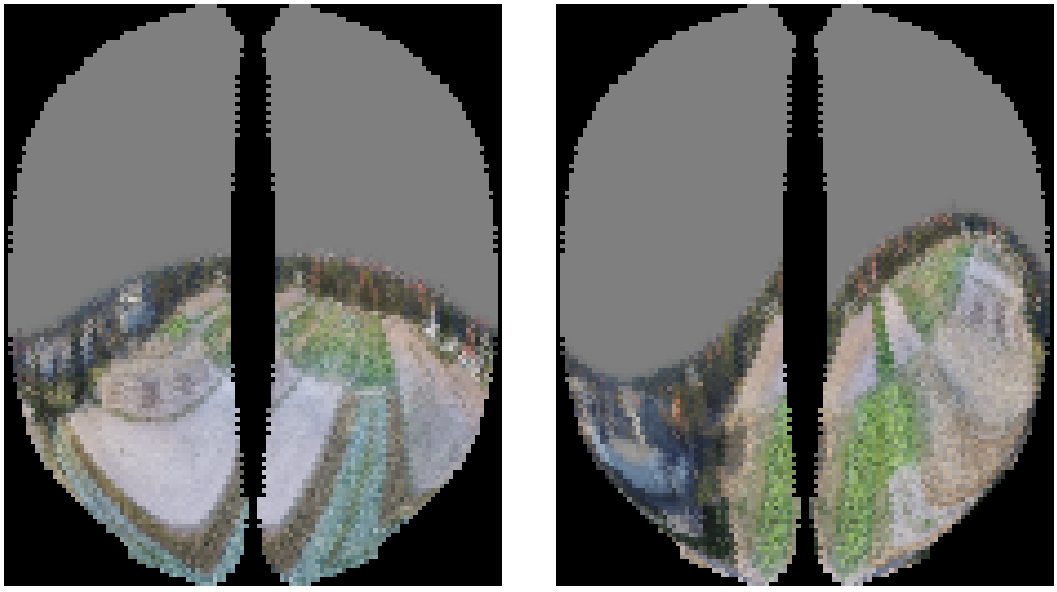

(a) Bee perspective during straight flight (b) Bee perspective during right turn

Figure 30: Comparison of the field's shape during straight flights and right turns in the perspective of a honeybee. Colours of pixels near the horizon often change due to rotations and alternations between showing a part of the field and the grey background.

Nevertheless, the results of the autoencoder approach would certainly be the same with appropriate textures with ultraviolet light because they are determined by different shapes of the field in bee view images instead of the field's texture.

The second approach for correlating a bee's vision with its brain activity was to project spike rates onto the environment in the range of the bee's viewing area. Figure 25 and Figure 26 show regions with high spike rates in turns in most rounds with exception of the last two flights on 3rd September, 2018. These are results when the average spike rate assigned to each pixel on the map is calculated. However, the maximum spike rate values per pixel reveals partly other regions with high spike rates as shown for the first flight of 3rd September, 2018 in Figure 27. Maps of rounds with maximum spike rates per pixel have such regions on the route before, over and after flying over the feeder on the first flight line. This is not visible as clearly or at all when the average spike rate per pixel is used.

The reason is the difference in the number of assigned spike rate values per pixel while flying the route as depicted in Figure 31a. Pixels in the region of the feeder count up to about 200 spike rate values because the copter flies three times over this part while pixels in rounds have only about 130 assigned.

Since the first part of the first line is flown at the start and end of a round, pixels in this location count more than 175 spike rate values. Regions further away from the flight route are assigned less than 50 spike rate values. Hence, high spike rate values near the feeder vanish when the average is taken. Regions further away from the route can show high spike rates despite being in the bee's viewing area only for a short time.

Since the SWCs of the spike rate as well as the autoencoder approach have already shown that turns are detectable in the spike rate, turns were excluded before projecting the spike rate onto the map in Figure 28. Nevertheless, regions with high spike rates on the field are not at the same locations across all rounds. That is why no unique location on the field could be found which always triggers higher spike rates. In particular, the feeder's location is not identifiable on the projected spike rate maps. Besides, the distribution of the amount of assigned spike rate values to pixels on the map is still not uniformly distributed as shown in Figure 31b.

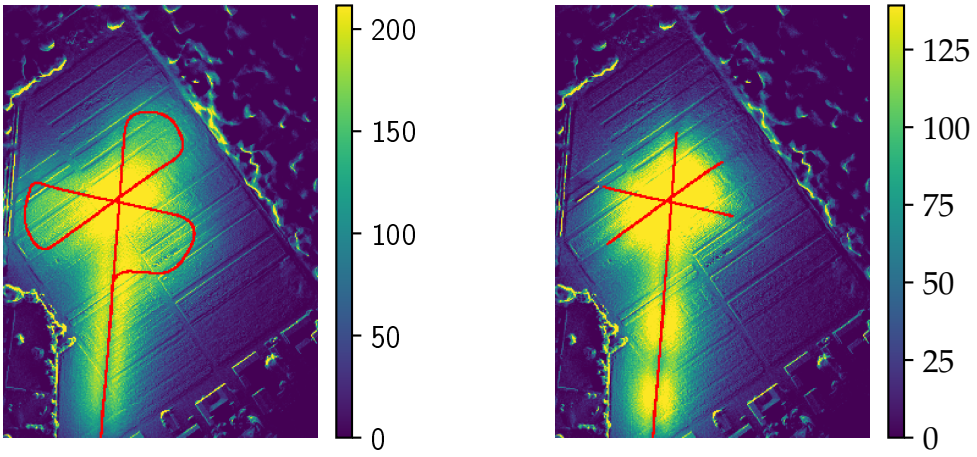

(a) One round

(b) Lines of one round

Figure 31: Distribution of viewing areas after one flight round. Maps show how many spike rate values are assigned to each pixel after the first round of the first flight in September during mapping. Map (a) depicts a complete round while in map (b) only straight lines are used as shown by the red flight trajectory.

Concluding, all results of the executed approaches show high brain activity of the honeybee when NeuroCopter flies turns. In fact, these animals respond to unintended rotations of their body in order to keep their field of vision stable [18]. That might be the reason for higher spike rate in turns. Moreover, none of the approaches showed evidence that training bees influenced their brain activity although the training let them know a feeder on the field before they flew over it with NeuroCopter. Results indicate that the feeder did not change the outcome of the experiment.

8 Conclusion

8.1 Results

The project NeuroCopter investigates navigational capabilities of honeybees. A quadcopter was modified in such a way that neural signals of an attached living bee can be measured while flying. The goal of this thesis was to find correlations between the visual perception of a bee with its brain activity in conducted experiments with NeuroCopter. Honeybees were trained to recognize a feeder on the experiment's field. Besides the feeder, the bee could use further landmarks such as paths for navigation. Since neurons in a region of the bee's brain were measured which capture visual information, correlations between brain activity and landmarks while flying were expected.

The copter flies a programmed route autonomously multiple times in experiments. That is why the visual perception of the bee is similar at each repetition. If vision and brain activity are correlated, the brain activity must also be similar across multiple iterations of the route. In fact, correlations between brain activities recorded in different rounds proved that this is the case.

Furthermore, vision and brain activity were correlated in two ways. First, a virtual 3D model of the experiment's field and its surrounding was created. Recorded telemetry data of the copter enabled accurate transfer of its location and tilt during flight to the simulation. In addition, images were rendered such that the field of view of a bee is represented. Next, an autoencoder processed the bee view images of the simulation. Thereby, encoded parts of these images were correlated with neural activity of the bee.

Another approach was to project the spike rates in the bee's visual range onto the field. Projecting changes of the bee's brain activity during the route creates a map that shows regions which triggered high neural activity. Both approaches, correlating processed images of a simulated flight as well as projecting brain activity to a map, revealed correlations between vision and neural activity in turns. However, no correlations with the feeder or other landmarks could be identified. Nevertheless, developed approaches in this thesis can be extended and applied to further experiments with NeuroCopter.

8.2 Outlook

This thesis analysed the first experiments of NeuroCopter which recorded brain activity of a living honeybee during flights. Naturally, there are several ways to improve further experiments.

On the one hand, hardware improvements are possible. A camera which can capture ultraviolet light is needed for textures on a virtual 3D model of the experiment's environment as described in Section 7. The current model does not reflect the colour range of honeybees. While this might not change the result of the approach which used the autoencoder described in Section 5, it could improve other methods which use the 3D model in the future. Another hardware improvement would be to reconstruct the holding of the bee such that the animal can move freely and navigate the copter. Thereby it would be clear that the bee does not counterbalance rotations of the copter which influences measured brain activity as described in Section 7.

On the other hand, the experiments can be improved. Flying with NeuroCopter in environments with more landmarks such as trees could result in notable correlations between vision and brain activity. The fields of the Julius Kühn-Institut contain only paths which could be used as landmarks for orientation. Moreover, different flight routes are needed to improve results of the projection of brain activity onto a map. As shown in Figure 31 the amount of assigned values to each pixel varies which influences the result. A flight route in form of a grid would be more suitable to prevent this problem. In addition, the grid route could be aligned such that parts are flown right over a path, other parallel to a path and some not aligned or crossing paths. This could enable analysation of correlations of brain activity with paths.

Generally, more experiment data are needed to compare results. Furthermore, instead of analysing brain activity in form of spike rates, interspike-intervals could be used.

Finally, the workgroup of the project NeuroCopter wants to improve the derivation of bees' neurons. This process is very laborious. The electrode needs to be prepared and inserted manually into the brain. It is important that the electrode is not moved too far in order to avoid damaging cells. As a result, the derivation is error-prone which leads to a high number of dead bees. Therefore, preparing a bee for using with NeuroCopter is time-consuming. Neural activity is displayed on an oscilloscope. The activity is triggered by moving hands or a torch in front of the bee's eyes. This movement is similar to the flown turns which correlate strongly as described in this thesis. However, there might be other neurons which are more important for the navigational processing of the bee and are not covered by this selection.

Bibliography

[1] M. Beekman and F. L. W. Ratnieks. Long-range foraging by the honeybee, apis mellifera l. *Functional Ecology*, 14(4):490–496, 2000.

[2] Philipp Breinlinger. Neurocopter: Optimierte verfahren zur positions- und lagebestimmung eines biomimetischen quadrocopters. *BioRobotics-Lab*, 2013.

[3] Alexis Buatois, Pichot Cécile, Patrick Schultheiss, Jean-Christophe Sandoz, Claudio Lazzari, Lars Chittka, Aurore Avargues-Weber, and Martin Giurfa. Associative visual learning by tethered bees in a controlled visual environment. *Scientific Reports*, 7, 12 2017.

[4] James F. Cheeseman, Craig D. Millar, Uwe Greggers, Konstantin Lehmann, Matthew D. M. Pawley, Charles R. Gallistel, Guy R. Warman, and Randolf Menzel. Reply to cheung et al.: The cognitive map hypothesis remains the best interpretation of the data in honeybee navigation. *Proceedings of the National Academy of Sciences*, 111(42):E4398–E4398, 2014.

[5] James F. Cheeseman, Craig D. Millar, Uwe Greggers, Konstantin Lehmann, Matthew D. M. Pawley, Charles R. Gallistel, Guy R. Warman, and Randolf Menzel. Way-finding in displaced clock-shifted bees proves bees use a cognitive map. *Proceedings of the National Academy of Sciences*, 111(24):8949–8954, 2014.

[6] Lin Chen, Shaowu Zhang, and Mandyam V. Srinivasan. Global perception in small brains: Topological pattern recognition in honey bees. *Proceedings of the National Academy of Sciences*, 100(11):6884–6889, 2003.

[7] Allen Cheung, Matthew Collett, Thomas S. Collett, Alex Dewar, Fred Dyer, Paul Graham, Michael Mangan, Ajay Narendra, Andrew Philippides, Wolfgang Stürzl, Barbara Webb, Antoine Wystrach, and Jochen Zeil. Still no convincing evidence for cognitive map use by honeybees. *Proceedings of the National Academy of Sciences*, 111(42):E4396–E4397, 2014.

[8] Lars Chittka and Jeremy Niven. Are bigger brains better? *Current Biology*, 19(21):R995 – R1008, 2009.

[9] Benjamin Daumenlang. Anbindung einer flugsteuerungsplatine an einen flugsimulator. *BioRobotics-Lab*, 2013.

[10] Jacqueline Degen, Thomas Hovestadt, Mona Storms, and Randolf Menzel. Exploratory behavior of re-orienting foragers differs from other flight patterns of honeybees. *PLOS ONE*, 13:e0202171, 08 2018.

[11] Jacqueline Degen, Andreas Kirbach, Lutz Reiter, Konstantin Lehmann, Philipp Norton, Mona Storms, Miriam Koblofsky, Sarah Winter, Petya B.

Georgieva, Hai Nguyen, Hayfe Chamkhi, Uwe Greggers, and Randolf Menzel. Exploratory behaviour of honeybees during orientation flights. *Animal Behaviour*, 102:45 – 57, 2015.

[12] Jacqueline Degen, Andreas Kirbach, Lutz Reiter, Konstantin Lehmann, Philipp Norton, Mona Storms, Miriam Koblofsky, Sarah Winter, Petya B. Georgieva, Hai Nguyen, Hayfe Chamkhi, Hanno Meyer, Pawan K. Singh, Gisela Manz, Uwe Greggers, and Randolf Menzel. Honeybees learn landscape features during exploratory orientation flights. *Current Biology*, 26(20):2800 – 2804, 2016.

[13] Adrian G. Dyer, Angelique C. Paulk, and David H. Reser. Colour processing in complex environments: insights from the visual system of bees. *Proceedings of the Royal Society B: Biological Sciences*, 278, 2010.

[14] Harald E. Esch, Shaowu Zhang, Mandyan V. Srinivasan, and Juergen Tautz. Honeybee dances communicate distances measured by optic flow. *Nature*, 411(6837):581, May 2001.

[15] Kunihiko Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36(4):193–202, Apr 1980.

[16] Sabine Gillner, Anja M. Weiß, and Hanspeter A. Mallot. Visual homing in the absence of feature-based landmark information. *Cognition*, 109(1):105 – 122, 2008.

[17] Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew M Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017.

[18] M. R. IBBOTSON. A motion-sensitive visual descending neurone in apis mellifera monitoring translatory flow-fields in the horizontal plane. *Journal of Experimental Biology*, 157(1):573–577, 1991.

[19] Michael Ibbotson, Yu-Shan Hung, H Meffin, Norbert Boeddeker, and M V. Srinivasan. Neural basis of forward flight control and landing in honeybees. *Scientific Reports*, 7, 12 2017.

[20] Tim Landgraf, Benjamin Wild, Tobias Ludwig, Philipp Nowak, Lovisa Helgadottir, Benjamin Daumenlang, Philipp Breinlinger, Martin Nawrot, and Raúl Rojas. Neurocopter: Neuromorphic computation of 6d egomotion of a quadcopter. In *Proceedings of the Second International Conference on Biomimetic and Biohybrid Systems*, Living Machines'13, pages 143–153, Berlin, Heidelberg, 2013. Springer-Verlag.

[21] Cambridge Electronic Design Limited. Spike2 software - Advanced fea-

tures. http://ced.co.uk/products/spkovaf, last accessed on 15th August, 2019.

[22] Martin Lindauer. Dauertänze im bienenstock und ihre beziehung zur sonnenbahn. *Naturwissenschaften*, 41(21):506–507, Jan 1954.

[23] Tobias Ludwig. Neurocopter - eine fliegende experimentierplattform zur erforschung der hirnaktivität von honigbienen. *BioRobotics-Lab*, 2016.

[24] Randolf Menzel, Uwe Greggers, Alan Smith, Sandra Berger, Robert Brandt, Sascha Brunke, Gesine Bundrock, Sandra Hülse, Tobias Plümpe, Frank Schaupp, Elke Schüttler, Silke Stach, Jan Stindt, Nicola Stollhoff, and Sebastian Watzl. Honey bees navigate according to a map-like spatial memory. *Proceedings of the National Academy of Sciences*, 102(8):3040–3045, 2005.

[25] Randolf Menzel, Andreas Kirbach, Wolf-Dieter Haass, Bernd Fischer, Jacqueline Fuchs, Miriam Koblofsky, Konstantin Lehmann, Lutz Reiter, Hanno Meyer, Hai Nguyen, Sarah Jones, Philipp Norton, and Uwe Greggers. A common frame of reference for learned and communicated vectors in honeybee navigation. *Current Biology*, 21(8):645 – 650, 2011.

[26] Randolf Menzel, Lea Tison, Johannes Fischer-Nakai, James Cheeseman, Maria Sol Balbuena, Xiuxian Chen, Tim Landgraf, Julian Petrasch, Johannes Polster, and Uwe Greggers. Guidance of navigating honeybees by learned elongated ground structures. *Frontiers in Behavioral Neuroscience*, 12:322, 2019.

[27] Marcel Mertes, Laura Dittmar, Martin Egelhaaf, and Norbert Boeddeker. Visual motion-sensitive neurons in the bumblebee brain convey information about landmarks during a navigational task. *Frontiers in Behavioral Neuroscience*, 8:335, 2014.

[28] Angelique C. Paulk, Andrew M. Dacks, James Phillips-Portillo, Jean-Marc Fellous, and Wulfila Gronenberg. Visual processing in the central bee brain. *Journal of Neuroscience*, 29(32):9987–9999, 2009.

[29] Julian N. G. Petrasch. A flying platform for behavioral and electrophysiological studies in honeybee navigation. *BioRobotics-Lab*, 2018.

[30] Jenny A. Plath, Brian V. Entler, Nicholas H. Kirkerud, Ulrike Schlegel, C. Giovanni Galizia, and Andrew B. Barron. Different roles for honey bee mushroom bodies and central complex in visual learning of colored lights in an aversive conditioning assay. *Frontiers in Behavioral Neuroscience*, 11:98, 2017.

[31] Johannes Polster. Simulating bee vision: Conceptualization, implemen-

tation, evaluation and application of a raycasting rendering engine for generating bee views. *BioRobotics-Lab*, 2017.

[32] Scott D Roth. Ray casting for modeling solids. *Computer Graphics and Image Processing*, 18(2):109 – 144, 1982.

[33] H. Sakoe and S. Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1):43–49, February 1978.

[34] Hiroaki Sakoe and Seibi Chiba. A dynamic programming approach to continuous speech recognition. In *Proceedings of the Seventh International Congress on Acoustics, Budapest*, volume 3, pages 65–69, Budapest, 1971. Akadémiai Kiadó.

[35] Rainer Schlittgen. *Einführung in die Statistik, Analyse und Modellierung von Daten*. De Gruyter, 12 edition, 2012.

[36] Reinhard Seidl and Walter Kaiser. Visual field size, binocular domain and the ommatidial array of the compound eyes in worker honey bees. *Journal of comparative physiology*, 143(1):17–26, Mar 1981.

[37] Sadia Shakil, Chin-Hui Lee, and Shella Dawn Keilholz. Evaluation of sliding window correlation performance for characterizing dynamic functional connectivity and brain states. *NeuroImage*, 133:111 – 128, 2016.

[38] Aung Si, Mandyam V. Srinivasan, and Shaowu Zhang. Honeybee navigation: properties of the visually driven 'odometer'. *Journal of Experimental Biology*, 206(8):1265–1273, 2003.

[39] Karin Steijven, Johannes Spaethe, Ingolf Steffan-Dewenter, and Stephan Härtel. Learning performance and brain structure of artificially-reared honey bees fed with different quantities of food. *PeerJ*, 5:e3858, October 2017.

[40] Benjamin Wild. Bestimmung der 6d-eigenbewegung eines fliegenden roboters anhand von monokularen messungen des optischen flusses. *BioRobotics-Lab*, 2014.

[41] Eric Zetzsche. Kollisionserkennung für single-/multicopter mit hilfe des optischen flusses. *BioRobotics-Lab*, 2014.

A Appendix

A.1 Code

The bee's spike rates were analysed and correlated with *Python*. The code was developed in *Jupyter notebooks*. The following repository contains notebooks to recreate all figures of this thesis and additional code used during the analysis:

https://github.com/tobiasschuelke/NeuroCopter-Analysis

A.2 Additional Figures

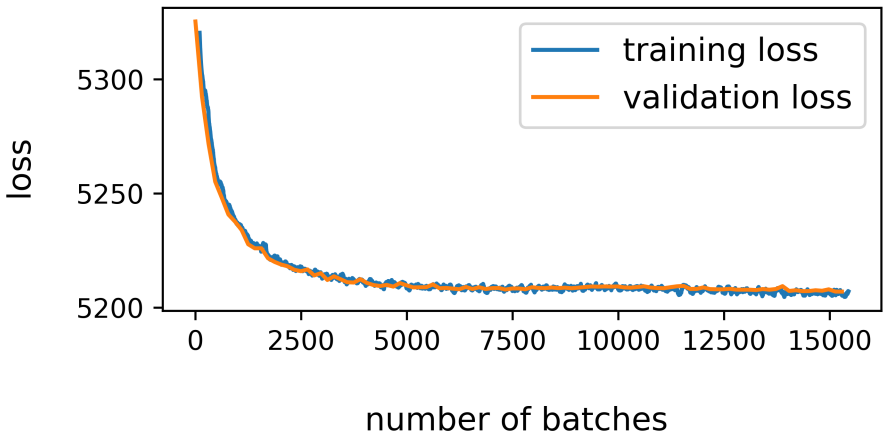

Figure 32: Training and validation loss of the autoencoder which was trained for 100 epochs. The autoencoder is used in Section 5. The training loss is smoothed by a rolling mean window of size 100 and the validation loss is depicted as the mean per epoch.

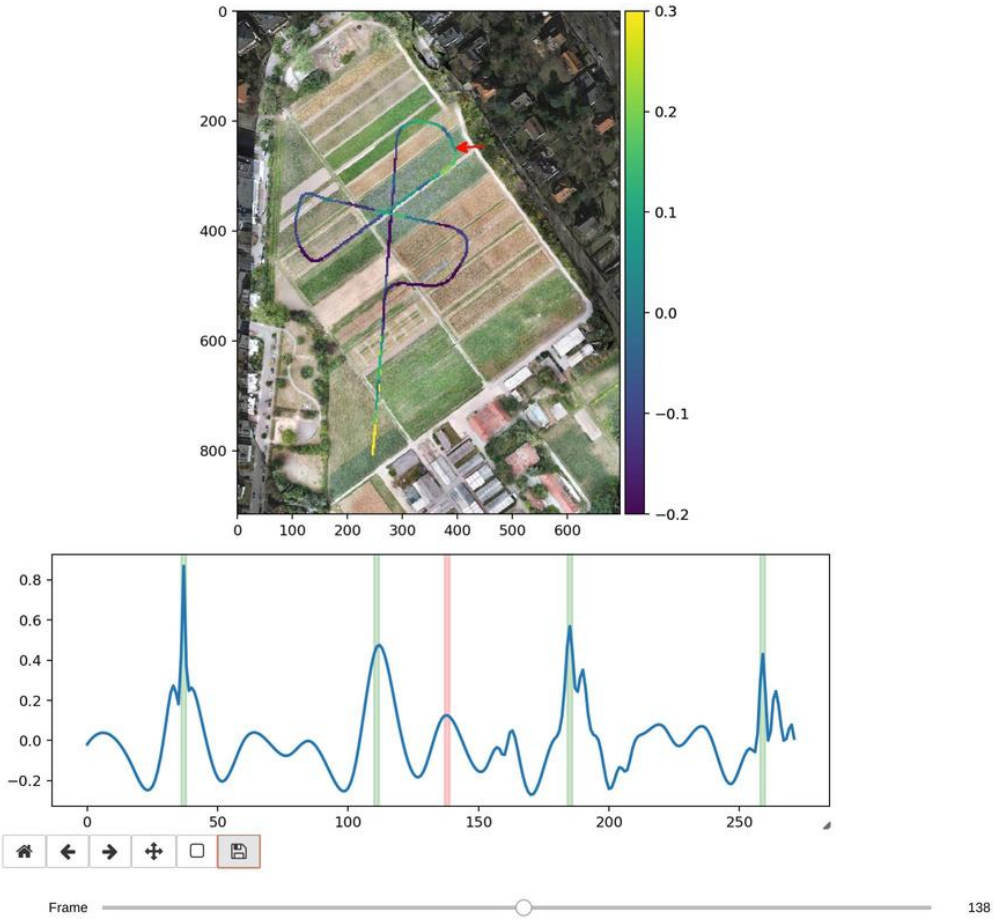

Figure 33: Overview to compare positions in the sliding window correlation graph with the corresponding position of the copter. The map displays the flight route of one round. Here, the first round is used for SWC. Colours along the route show correlation values and the arrow shows the location of the copter. The graph shows the SWC of the whole flight. Round starts are highlighted in green. The slider at the bottom is used to select positions in the graph. The current selected index of the SWC in the graph is marked in red. The slider is used to change the position in the graph and updates the location of the copter (arrow) in the map accordingly. The correlation values along the flight route are updated when the position in the graph is moved to another round.

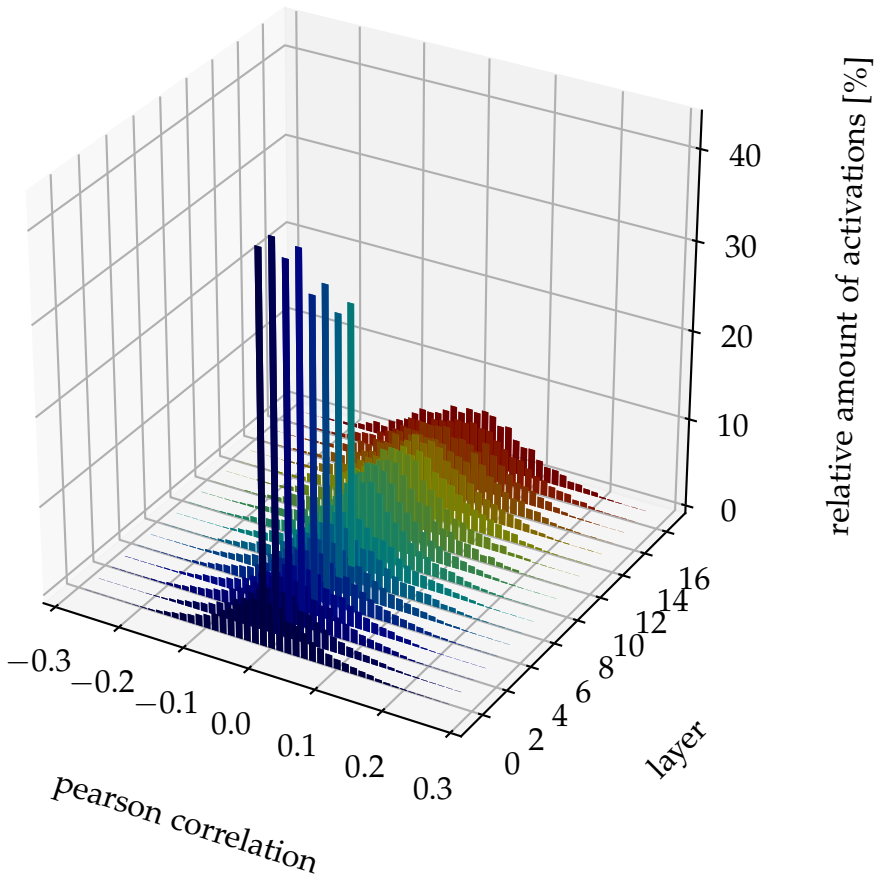

Figure 34: Distribution of all correlations of activations with spike rate during all rounds of the second flight on 3rd September, 2018. Correlations are separated by layers of the corresponding activation. Histograms of each layer are normalised to show the relative distribution of correlation values per layer.

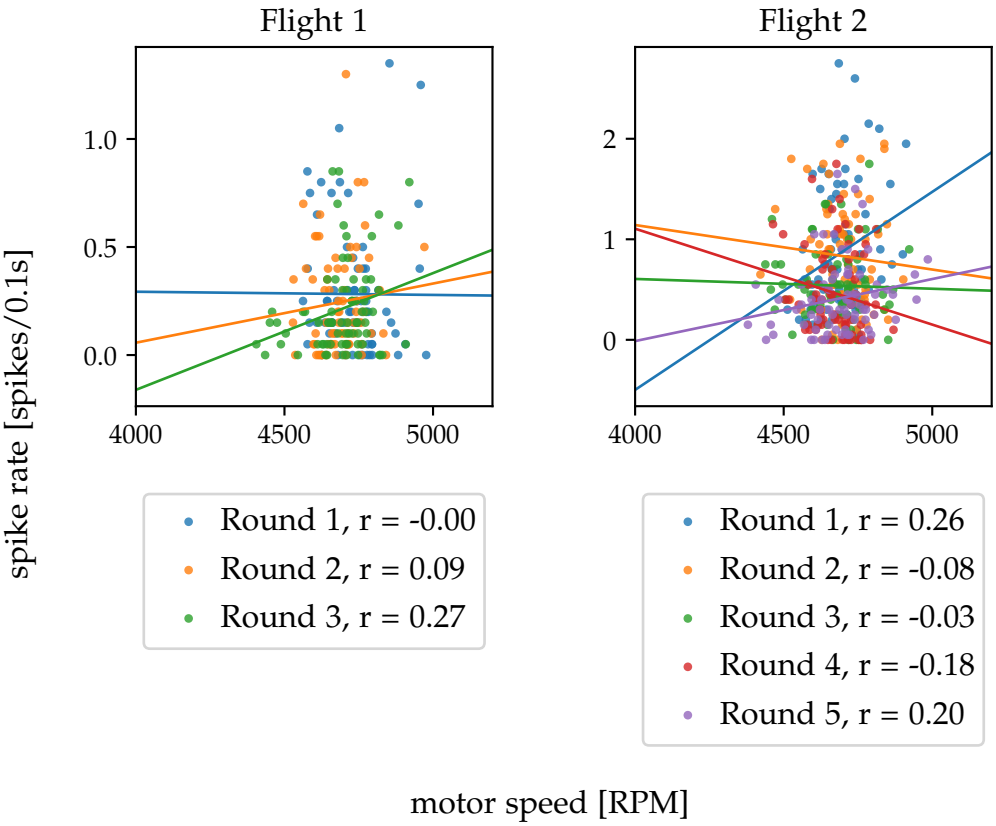

Figure 35: Scatter plot of spike rate and mean motor speed in August. Mean speed of the copter's four motors is used. The average of the spike rate and mean motor speed are calculated in two second sections per flight round. Averaged values are shown in scatter plots together with their linear regressions. Correlation values of the mean spike rate and motor speed in two second sections are annotated in the legends. The motor speed is measured in rotations per minute.

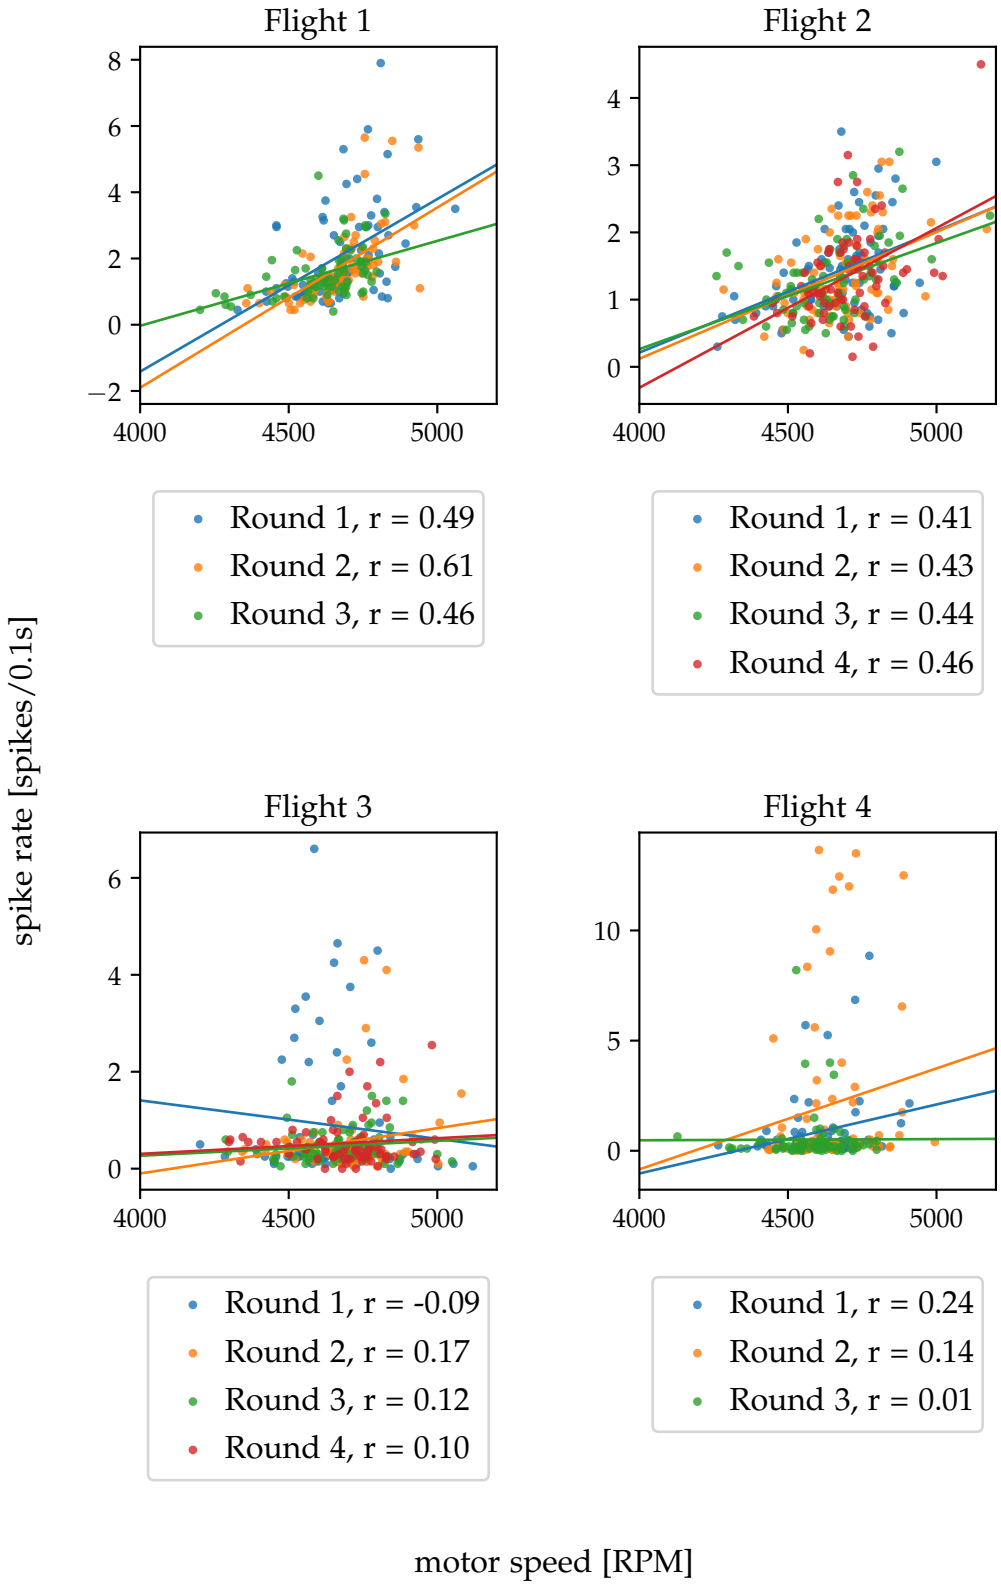

Figure 35: Scatter plot of spike rate and mean motor speed in September (continued).

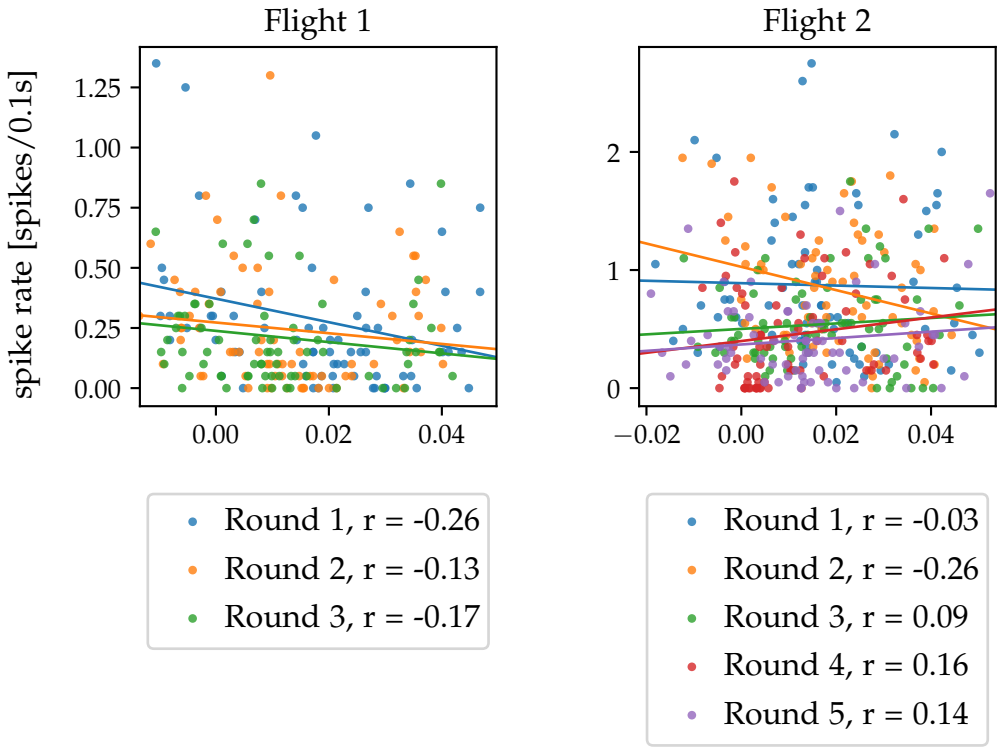

Figure 36: Scatter plot of spike rate and mean acceleration in August. Mean of the acceleration along the x, y and z axes is used. The average of the spike rate and mean acceleration are calculated in two second sections per flight round. Averaged values are shown in scatter plots together with their linear regressions. Correlation values of the mean spike rate and acceleration in two second sections are annotated in the legends. The acceleration is measured in metres per second.

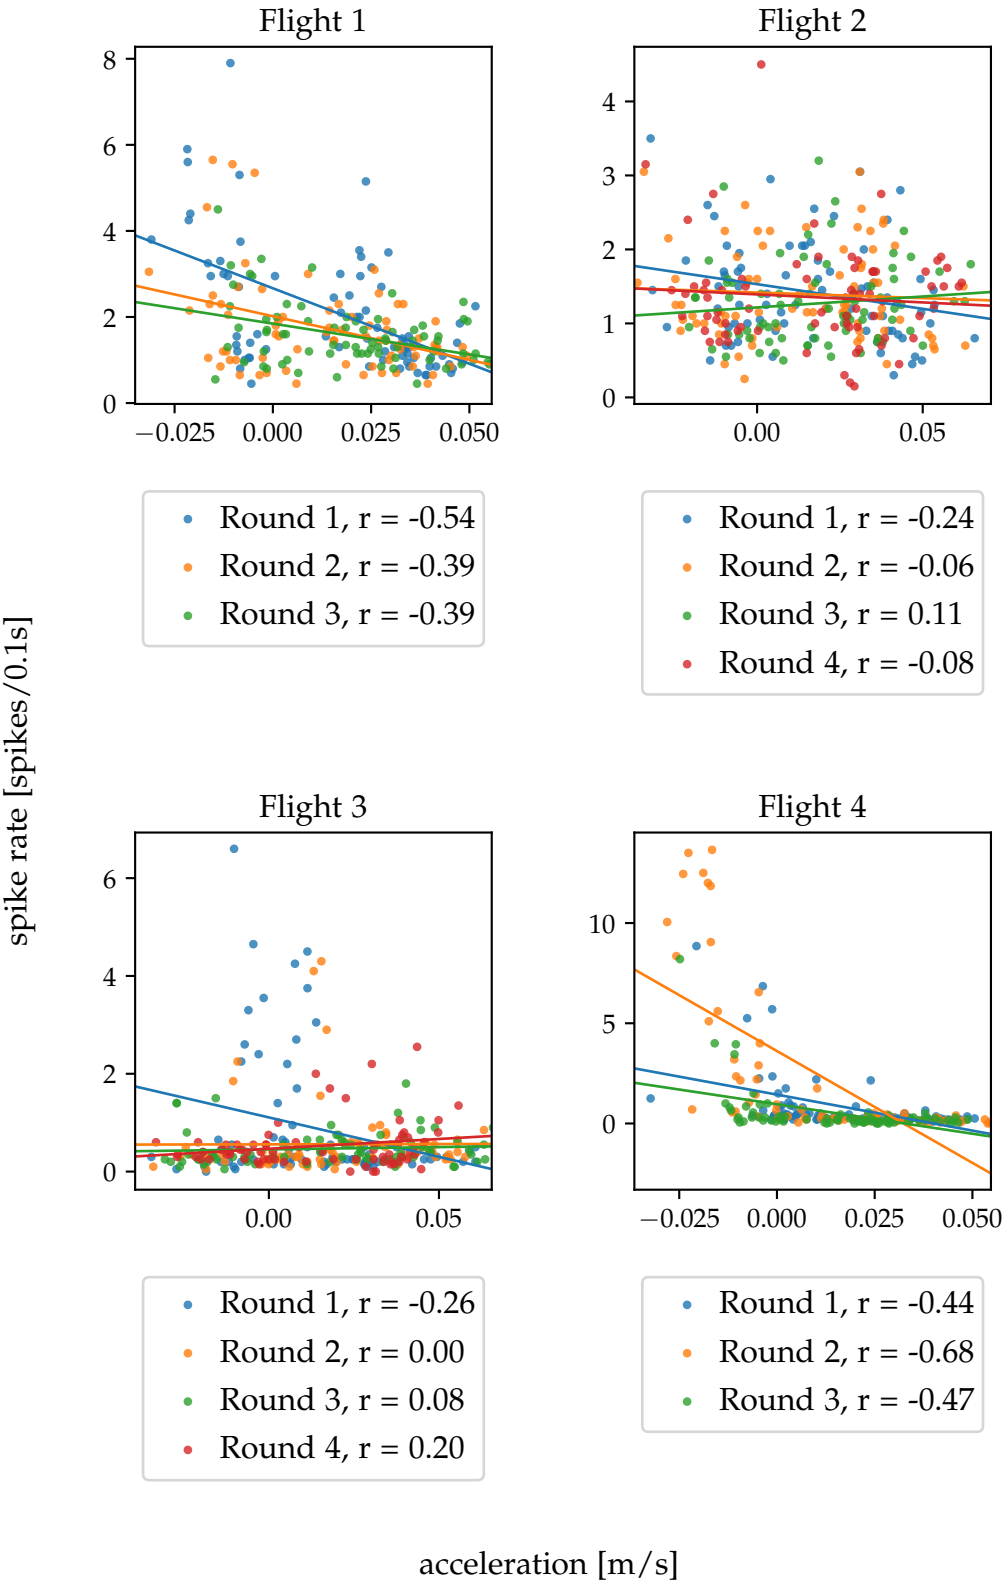

Figure 36: Scatter plot of spike rate and mean acceleration in September (continued).