Identifying All Individuals in a Honeybee Hive – Progress Towards Mapping All Social Interactions

Cris L. Luengo Hendriks[1], Ziquan Yu[1], Antoine Lecocq[2,3], Teatske Bakker[2], Barbara Locke[2], Olle Terenius[2]
[1]*Centre for Image Analysis, Swedish University of Agricultural Sciences & Uppsala University*
[2]*Department of Ecology, Swedish University of Agricultural Sciences*
[3]*Currently at: Department of Agriculture & Ecology, University of Copenhagen*
cris@cb.uu.se

Abstract

Here we present work in progress towards a fully automatic system that monitors a honeybee hive over many days, deriving information on the position and velocity of each bee, and detecting and identifying each instance of a social interaction. Each bee is tagged with a unique identifier, enabling the system to know exactly which individuals interacted in each case. The final result should be a map of all interactions, from which it is possible to derive, for example, a sociogram.

1. Introduction

The honeybee is a social animal, and a bee hive is a complete, closed society. Within this society, much like a human society, some individuals physically interact with many different individuals, while others have only few connections. This makes a hive an ideal model system to study social-contact–related processes such as disease transmission.

Advantages of studying a model system rather than a group of humans are numerous. The most obvious one is the ability to directly interfere with the model system (e.g. by infecting some individuals), in ways one cannot do with a group of human test subjects. Information on social interactions in groups of humans must be obtained by questionnaires and the like, which leaves the possibility of missing important interactions that the people involved in are not willing to disclose; in contrast, it is possible to directly observe a bee hive and detect all interactions. Additionally, interactions outside the studied human test group cannot be accounted for, but might influence the system; a bee hive is a closed system, interactions outside of it are rare.

Our goal is to observe a small hive, with around 1000 individuals, over the course of up to 10 days, in an environment that is as natural to the bees as possible. Within this hive, each individual has a unique tag on its thorax, that allows the system to identify the individual and assign interactions unequivocally. We currently have many of the subsystems ready, and have preliminary results showing our ability to identify and track bees.

1.1 Current state of the art

Automatic video analysis of bee behaviour has been done before, but at a much more restricted scale. There have been studies of the wiggle dance, in which the dancing bee was automatically tracked and the dance analysed [1]. This involves tracking a single bee, manually picked, over the course of minutes. Another paper presents a method to track many bees in a hive [2]. This method is based on vector field quantization and a rough segmentation of the field of view to try to detect as many individual bees as possible. A subset of the hive can be tracked over minutes, but identification of individuals is not possible.

2. Materials and methods

2.1. The observation hive

We use a standard observation hive (w × h × d = 43.5 × 52.5 × 5.5 cm), in which two standard hive frames (37 × 22 cm, see Figure 1) fit vertically one above the other. A plexiglass sheet on either side encloses the hive. The observation hive is installed in a small, windowless room, and a tube connects the hive to the outside, so that bees can go out and forage.

The setup allows filming of one side or both sides of the frames at the same time. If it is deemed relevant for the experiment to observe bees on both sides of the frame, it is possible to restrict the bees to the bottom half of the hive. Otherwise we can restrict the bees to one side of the hive. In either case, the surface area to be filmed is the same (two frame sides).

We sprayed a thin film of fluon on the inside of the plexiglass sheet using an airbrush. Fluon is slippery, and avoids, to a large extent, bees walking on the glass. If the film is thin enough, it is fully transparent.

2.2. The camera and illumination system

To minimize the impact of observation on the bee behaviour, we keep the room dark, illuminating the bees with near infra-red light (850 nm). According to the literature, bees are not sensitive to infra-red light [3]. LED lights at 850 nm are inexpensive and easy to obtain. We have eight 250 mW units. The problem is to illuminate the honeycomb behind a plexiglass sheet, which is highly reflective. We have build a light diffuser, but find that the most effective way of illuminating the field of view is through large angles, such that the reflected light does not reflect into the camera. However, both these methods reduce the total amount of light reaching the camera.

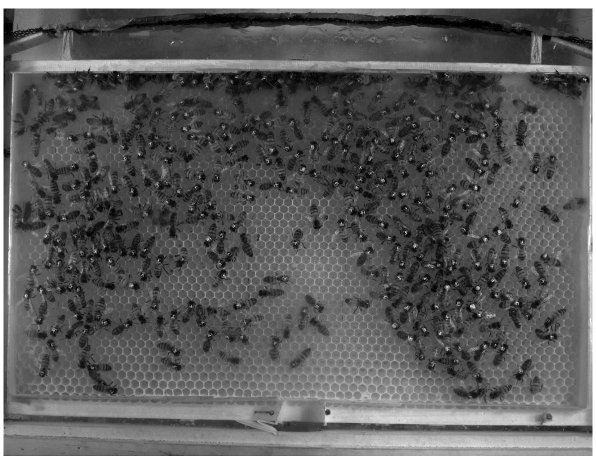

Figure 1: One bee hive frame captured at high resolution with 850 nm illumination.

We have two Basler Scout scA1600-14gm cameras, with Fujinon HF16HA-1B lenses and 850 nm bandpass filters. These filters block out all light not coming from the illumination system, and are meant to protect the observation system from light caused by, e.g. opening the door to the room where the hive is kept. As discussed later, reading the tags requires a constant, controlled illumination.

The cameras are placed 80 cm from the honeycomb, such that each camera's field of view completely covers one of the frames (Figure 1), and aligning their optical axis with the normal of the observed surface. We can place the two cameras to observe both sides of one frame, or one side of two frames stacked vertically, depending on the requirements of the experiment. Observing only one side of the hive is simpler, because the illumination is easier to set up properly.

The cameras record 14 monochrome frames per second, each frame has 1628 × 1236 pixels. The resolution is good enough to read tags on the bee's thorax (Figure 2). Currently we record the video feed and process it offline, but we hope to ultimately process the video in real time. It is important to use a lossless compression method, as compression artefacts can make it impossible to correctly identify the tags.

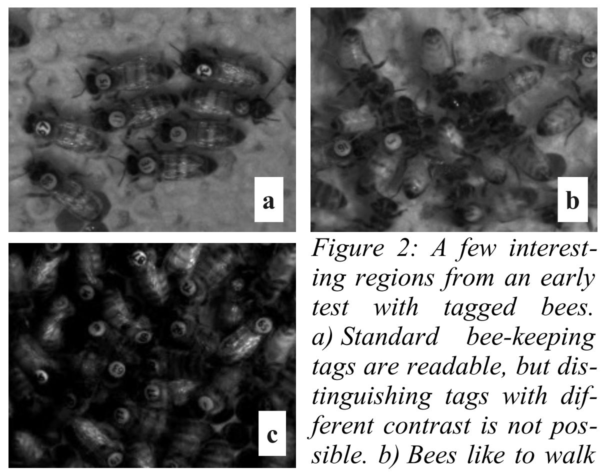

Figure 2: A few interesting regions from an early test with tagged bees. a) Standard bee-keeping tags are readable, but distinguishing tags with different contrast is not possible. b) Bees like to walk on the plexiglass cover. c) When bees are hugging close together, it is difficult to see which parts of the bees belong to each of the tags.

2.3. The tags

We first intended to use the tags commonly used in bee-keeping to mark the queen of each hive. These are small, circular, plastic tags (3 mm diameter) that can be glued onto the thorax of a bee (Figure 2). They come in different colours, and are marked with a number between 1 and 99. Since we decided against using colour video, our options for extending the range of numbers was limited to tags with different contrast (e.g. dark tags with white digits and white tags with black digits), and adding a mark on the abdomen (e.g. different lines in paint with a thin brush). This proved more difficult than we originally presumed. Furthermore, automatically identifying the number on the tag was not a trivial task either due to the low resolution, changing perspective as the bee moves around the hive, and specular reflection off the curved plastic surface (see Figure 2).

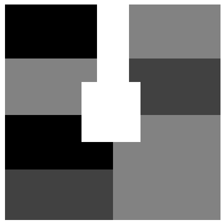

Figure 3: The custom tag design. 8 rectangular areas encode one of more than 6000 IDs using only three different grey values. The central white square is easy to detect, and the white line extending from it identifies the bee's orientation.

Figure 4: Two bees tagged and ready to be inserted into the observation hive.

Thus, we set out to design our own tags. We settled for a 3 × 3 mm square design as in Figure 3, which we could print on paper and easily cut. Designing our own tags allowed us to include more than only an ID: tags have a white dot in the middle that is easy to detect in software. The software thus only needs to detect these dots, rather than whole bees. The white dot extends into a line to one side of the square. This line indicates the direction towards the head of the bee. The rest of the surface contains areas of different grey values, encoding a unique ID for a bee (see Figure 3 and 4).

To encode 1000 unique numbers with two grey values would require ten areas. But, at the resolution with which we image the bees, we estimated that eight such areas is the maximum. Thus, we need to use three different grey values, giving a total of $3^8 = 6\,561$ combinations. Having six times as many combinations as needed is useful to encode an error correction mechanism, making the identification more robust. The eight areas that encode the ID are distributed four on each side of the tag. To determine which grey values are best distinguished under 850 nm light, we printed a uniform gradient from black to white, and photographed it with our camera system. In the resulting image, we looked for areas with a 33% and 66% intensity. The grey values used to print these areas (26% and 51% grey in our case) were the ideal ones to use, together with pure black and the pure white of the central dot, to generate the tags. Note that these values need to be calibrated for each specific combination of printer, ink, and illumination.

2.4. The software

We currently have code that detects tags in individual frames, tracks the tags across many frames, and reads the ID on the tags. Our detection algorithm simply detects small, bright dots with a combination of the Laplace of Gaussian filter, an erosion and a simple threshold. Bright points found that do not belong to a tag are discarded in a second step where we identify the orientation of the tag: there needs to be a white line extending from the central dot.

The tracking algorithm is based on the mean-shift procedure, but includes certain physical constraints, given by the direction of the bee and the maximum possible speed of a bee. We detect tags globally every 5 frames; this discovers bees missed in an earlier frame and bees entering the hive, as well as corrects for any errors in tracking. The global detection step takes more time than tracking all bees from one frame to the next, and therefore we need to limit how often this step is performed. The tracking is solely meant for speeding up the process.

To read the ID on the detected tags, we extract pixel values (with interpolation) along four lines perpendicular to the white line on the tag (yellow lines in Figure 5), as well as along two lines parallel to it (blue lines in Figure 5; these lines are actually not perfectly parallel so as to optimize our ability to estimate grey values correctly). The ID is determined independently from these two sets of values, and compared for consistency. Note that, because we can track bees across many frames, it is possible to do this procedure at several time steps along one track, increasing robustness of the identification.

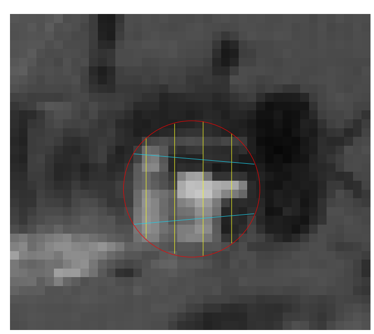

Figure 5: Grey values of the code on the tag are measured independently along two sets of lines.

Because up to now the system recorded the video, which we processed offline, no effort into speeding up these algorithms has been made. Consequently, it takes several days to process data for a few hours of video. The main problem with computational speed is caused by the huge amount of data collected by one camera. A computer with multiple high-end graphics cards is one possible solution [4].

Detection and identification of interaction has not yet been implemented. We currently simply record instances where two bees are close enough together and with the heads towards each other. Eventually, we will build a library of examples of interaction, which we can use to train a classification system that can identify the various types of interaction.

3. First results

In a very early experiment last year, we recorded several hours of video of 700 bees marked with the standard round numbered tags. These videos were created to study the effect of light on bees. We turned a light on and off during recording; this light was not visible in the video due to the IR bandpass filter on the camera. We showed a clear increase of activity (as measured by the increase of average speed of the bees) for a short time after turning on or off the light [5].

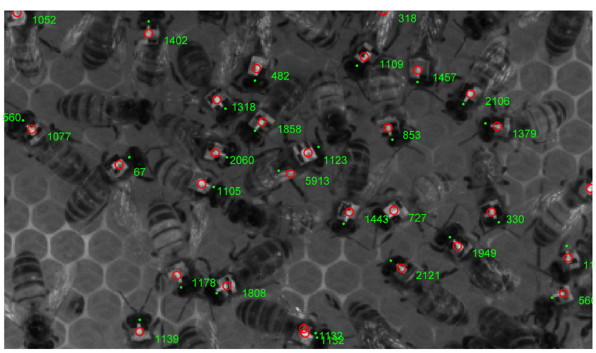

Figure 6: A portion of a video frame with red markers at detected tags, green dots indicating the head direction, and the tag ID. Note a specular reflection off a wing that is identified as tag #5913.

More recently, we recorded around 350 bees, tagged with our custom tags, with a single camera during seven days (producing around 1Tb of video a day). One frame is shown in Figure 1. The process taught us several important lessons, most notably that tags need to be rugged and well attached, as the wear is significant; we even saw bees helping each other rip the tag off. We applied the detection and tracking software on a short 30 s clip, in which 284 bees still had a tag (see Figure 6). It correctly detected 89% of the tags on average, but also detected an average of 11 dots (mostly glossy reflections) as tags. Compared to existing systems, this is a very high correct detection ratio, though we think these numbers can still be improved.

4. Conclusions and future work

Observing 1000 bees in a comfortable environment, fully automatically, is no easy task. We have developed a hive, illumination and camera system that can record bees in the dark. Standard tags used in beekeeping are not suitable for computer reading for several reasons: their surface is curved and shiny, causing specular reflections; the numbers can be hard to read when a tag is not perfectly vertical; and they do not have enough variation to distinguish 1000 individuals. We have developed a custom tag that avoids all these problems. Furthermore, detecting the tag is easier and we have added a feature that allows us to determine the orientation of the bee from the tag alone. This means that we do not need to segment or identify the actual bees, only their tags.

We have also developed some software to detect, track and identify the tags on the bees. This various algorithms need to be integrated into a user-friendly program, and they need to be optimized to handle the large amount of data in real time. We also will develop a set of algorithms to detect and identify different types of interaction between bees. Once all these software components are in place, we will be able to derive sociograms for the whole hive fully automatically. This will enable many interesting experiments to be performed.

References

[1] A. Veeraraghavan, R. Chellappa and M. Srinivasan, Shape-and-Behavior Encoded Tracking of Bee Dances. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(3):463–476, 2008.

[2] T. Kimura, M. Ohashi, R. Okada and H. Ikeno, A new approach for the simultaneous tracking of multiple honeybees for analysis of hive behavior. *Apidologie*, 42(5):607–617, 2011.

[3] K. von Frisch, *Bees: Their Vision, Chemical Senses, and Language*. Cornell University Press, Ithaca, New York, 1963.

[4] K.J. Batenburg et al., *FASTRA II: the world's most powerful desktop supercomputer*. http://fastra2.ua.ac.be/ (retrieved 2012-07-29).

[5] A. Lecocq, *The development of tool for network analysis in the honey bee (Apis mellifera L.)*. M.Sc. Thesis, Faculty of Natural Resources and Agricultural Sciences, Swedish University of Agricultural Sciences, 2011.