Social networks predict the life and death of honey bees

Benjamin Wild[1†*], David M Dormagen[1†], Adrian Zachariae[2†], Michael L Smith[4,5,6], Kirsten S Traynor[7], Dirk Brockmann[2,3], Iain D Couzin[4,5,6], Tim Landgraf[1*]

[†] equal contribution
*corresponding authors (b.w@fu-berlin.de, tim.landgraf@fu-berlin.de)
[1] Department of Mathematics and Computer Science, Freie Universität Berlin, Berlin, Germany
[2] Robert Koch Institute, Berlin, Germany
[3] Institute for Theoretical Biology, Humboldt University Berlin, Berlin, Germany
[4] Department of Collective Behavior, Max Planck Institute of Animal Behavior, Konstanz, Germany
[5] Centre for the Advanced Study of Collective Behavior, University of Konstanz, Konstanz, Germany
[6] Department of Biology, University of Konstanz, Konstanz, Germany
[7] Global Biosocial Complexity Initiative, Arizona State University, Tempe, USA

Abstract

In many social systems, an individual's role is reflected by its interactions with other members of the group (Gordon 2010, Pinter-Wollmann et al. 2014, Krause 2015, Farine & Whitehead 2015). In honey bee colonies (*Apis mellifera*), workers generally perform different tasks as they age, yet there is high behavioral variation in same-aged bees (Seeley 1982, Robinson 1992, Huang and Robinson 1996, Johnson 2010). It is unknown how social interactions within the colony relate to an individual's tasks throughout her life. We propose a new method to extract a single number from each individual's interaction patterns in multimodal social networks that captures her current role in the colony. This "network age" is better than biological age at predicting task allocation (+99%), survival (+157%), and activity patterns (+44-108%) and even predicts task allocation up to one week (around a sixth of her typical lifespan) into the future. Network age identifies distinct developmental paths and task changes throughout a bee's life: We show that individuals change tasks gradually and exhibit high task repeatability, and that same aged bees form stable behavioral subgroups in which they predominantly interact with one another. While we derived interaction networks by automatically tracking a fully tagged colony, we show that tracking only 5% of the bees is sufficient to extract a meaningful representation of the individuals' interaction patterns, demonstrating the feasibility of our method for detecting complex social structures with reduced experimental effort. Since network age more accurately predicts task allocation than biological age, it could be used in experimental manipulations to quantify shifts in the timing of task transitions as a response. We extend our method to extract interaction patterns relevant to other attributes of the individuals, such as their mortality, opening up a broad range of possible applications. Our approach is a scalable instrument to study individual behavior through the lens of social interactions over time in honey bees and other complex social systems.

Introduction

In complex systems, intricate global behaviors emerge from the dynamics of interacting parts. Within animal groups, studying the different ways the individuals can interact helps elucidate their function (Gordon 2010, Pinter-Wollmann et al. 2014, Krause 2015, Farine & Whitehead 2015). Social interaction networks have been used to investigate e.g. pair-bonding (Psorakis et al. 2012), inter-group "brokering" (Lusseau & Newman 2004), offspring survival (Cheney et al., 2016), cultural spread (Aplin et al. 2015), policing behavior (Flack et al. 2006), leadership in mobile groups (Strandburg-Peshkin et al. 2018), and the propagation of alarm response

(Rosenthal et al., 2015; Sosna et al 2019). As our ability to collect detailed social network data increases, so too does our need to develop tools for understanding the significance and functional consequences of these networks (Gomez-Marin et al. 2014). However, the task of identifying an individual's role in a group based on its characteristic patterns of interaction remains challenging, particularly with large numbers of individuals and multiple modes of interaction.

Honey bees are an ideal model system to study the relationship between social interactions and individual roles (Schneider & Lewis 2004, Naug 2008, Baracchi & Cini 2014). Task allocation in honey bees is characterized by temporal polyethism, i.e. bees gradually change the types of tasks they perform as they age - young bees care for brood in the center of the nest; old bees of around 15-20 days begin foraging outside of the nest (Seeley 1982, Huang & Robinson 1996). Due to the spatial organization of honey bee colonies, this change in task also results in a change of location (Seeley 1982). While honey bees are well-known for their elaborate social signals (e.g. waggle dance, shaking signal, stop signal; Frisch 1967, Nieh 2010, Seeley et al. 2012), they also exchange information through food exchange, antennation, or simple colocalization (Goyret & Farina 2003, Balbuena et al. 2012).

Despite extensive work on the social physiology of honey bee colonies (Seeley 1995), few works have studied interaction networks from a colony-wide or temporal perspective (Gernat et al. 2018, Hasenjager et al. 2020). While there is considerable variance in task allocation, even among bees of the same age, it is unknown how this variation is reflected in the social networks. Do we find specific patterns of interactions within groups of bees with similar tasks and how do these patterns change over time? Scaling up to the colony-level, can we extract meaning from the multitude of interactions and interaction types, among thousands of individuals in a honey bee colony?

What is “network age”?

To address these questions, we require methods that will track the tasks and social interactions of many individuals over consecutive days. We video recorded a full colony of individually marked honey bees (*A. mellifera*) at 3 Hz for 25 days (from 2016-08-01 to 2016-08-25) and obtained continuous trajectories for all individuals in the hive (Wario et al. 2015, Boenisch et al 2018). We used a two-sided single-frame observation hive with a tagged queen and started introducing individually tagged bees into the colony approximately one month before the beginning of the focal period (2016-06-28 to 2016-08-23). To make sure that no unmarked individuals emerged inside the hive, we replaced the nest substrate regularly (approx. every 21 days). In total, we recorded 1,921 individuals aged from 0 days to 8 weeks.

A worker’s task and the proportion of time she spends in specific nest areas are tightly coupled in honey bees (Seeley 1982). We annotated nest areas associated with specific tasks (e.g. brood area or honey storage) for each day separately (see S004), as they can vary in size and location over time (Smith et al. 2016). We then use the proportion of time an individual spends in these areas throughout a day as an estimate of her current tasks.

We calculated daily aggregated interaction networks from contact frequency, food exchange (trophallaxis), distance, and changes in movement speed after contacts (see S003). These networks contain all the pairwise interactions between all individuals over time. For each day and interaction type, we extract a compact representation that groups bees together with similar interaction patterns, using spectral decomposition (Belkin & Niyogi 2003, Luxburg 2007). We then combine each bees' daily representations of all interaction types and map them to a scalar value (network age) that best reflects the fraction of time spent in the task-associated areas using CCA (canonical-correlation analysis; Hotelling 1936, Knapp 1978). Note that network age only contains information from the social networks; the spatial data only guides *which* information to extract from the social networks. Network age thus compresses millions of data points per individual and day (1,920 potential interaction partners, each detected 127,501 ± 50,340 times on average per day, with four different interaction types) into a single number that represents each bee's daily position in the multimodal temporal interaction network. Since CCA is applied over the entire focal period, network age can only represent interaction patterns that are consistent over time. See Figure 1 for an overview and S005 for a detailed description of the methods.

Network age is a unitless descriptor. We scale it such that 90% of the values are between 0 and 40 to make it intuitively comparable to a typical lifespan of a worker bee during summer, and because biological age is commonly associated with task allocation in honey bees. This scaling can be omitted for systems where behavior is not coupled with biological age.

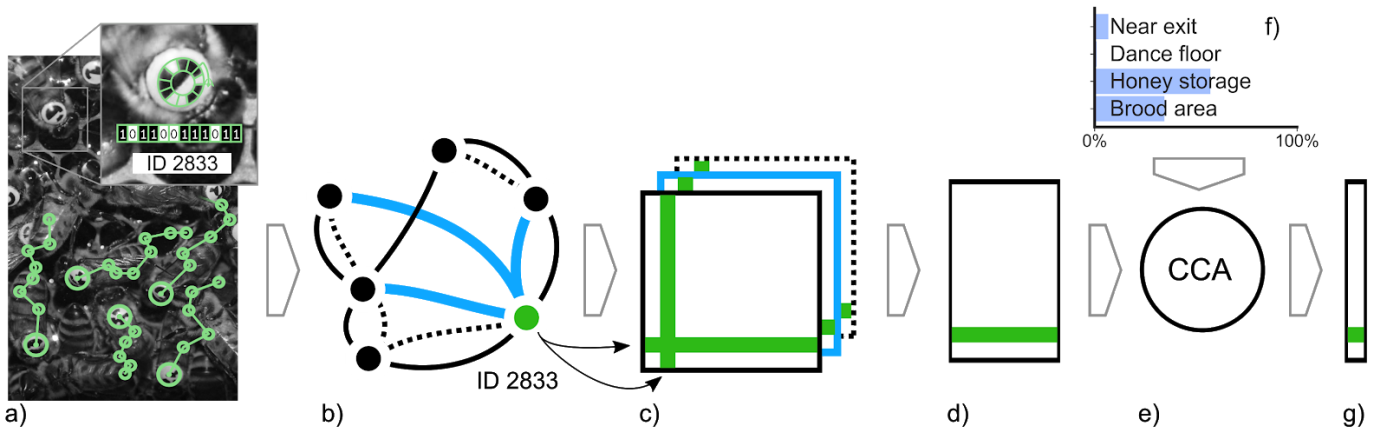

FIG 1: Network age, a one-dimensional descriptor of an individual's role within the colony, based on an individual's interaction pattern. Using the BeesBook automated tracking system, we obtain lifetime tracking data for individuals (a). These tracks are used to construct multiple weighted social interaction networks (b). We aggregate daily networks (c) to then extract embeddings that group bees together with similar interaction patterns, using spectral decomposition (d). Finally, we use a linear transformation (e; CCA = canonical-correlation analysis) that maximizes correlation with location preference (f) to compress them into a single number called "network age" (g).

Network age correctly identifies task allocation

We quantify to what extent network age, and therefore the social interaction patterns, reflect task allocation by using multinomial regression to predict the fraction of time each bee spends in the annotated nest areas (see S104). We then use McFadden's pseudo R^2 scores R^2_{McF} (McFadden 1973) to compare the regression performance of models based on an individual's biological age versus different variants of network age. Network age is twice as good as biological age at predicting the individuals' location preferences, and therefore their tasks (network age: median $R^2_{McF} = 0.682$, 95% CI [0.678, 0.687]; biological age: median $R^2_{McF} = 0.342$, 95% CI [0.335, 0.349]; 95% CI of effect size [0.332, 0.348], N=128; likelihood-ratio χ^2 test p << 0.001, N=26'403, table in S104, see S106 for details). Network age provides a better separability of time spent in task-associated nest areas than biological age (Fig 2a, example cohort in Fig 2c, S008 for all cohorts). Network age correlates with location because of the inherent coupling between tasks and nest areas. Still, it is not a direct measure of location: bees with the same network age can exhibit different spatial distributions and need not directly interact (see S103).

While we can improve the predictive power of network age by extracting a multi-dimensional descriptor instead of a single value (see S005 and S104 for details), the improvements for additional dimensions are marginal. This implies that a one-dimensional descriptor captures most of the information from the social networks that is relevant to the individuals' location preferences and therefore their tasks.

We experimentally demonstrated that network age robustly captures an individual's task by setting up sucrose feeders and identifying workers that foraged at the feeders (known foragers, N=40, methods in S105). We then compared the biological ages of these known foragers to their network ages. As expected, foragers exhibited a high biological age and a high network age, whereas biological age exhibited significantly larger variance than network age (Fig 2b; Levene's test (Levene 1960), performed on z-transformed values: p << 0.001, N=200). Indeed, while we observed a forager as young as 12 days old, that individual had a network age of 25.5, demonstrating that network age more accurately reflected her task than her biological age (Z-transformed values: biological age -0.46; network age 0.61).

Tagging an entire honey bee colony is laborious. However, by sampling subsets of bees, we find that network age is still a viable metric, even when only a small proportion of individuals are tagged and tracked. With only 1% of the bees tracked, network age is still a good predictor of task (median $R^2_{McF} = 0.516$, 95% CI [0.135, 0.705], N=128) while increasing the number of tracked individuals to 5% of the colony results in a R^2_{McF} value comparable to the completely tracked colony (5% of colony tracked: median $R^2_{McF} = 0.650$, 95% CI [0.578, 0.705], N=128; whole colony tracked: median $R^2_{McF} = 0.682$, 95% CI [0.678, 0.687], N=128; see S102). Similarly, we find that an approximation of network age can be calculated without annotated nest areas: Network age can be extracted in an unsupervised manner using PCA on the spectral embeddings of the different interaction type matrices (median $R^2_{McF} = 0.646$, 95% CI [0.641, 0.650], N=128, see S006).

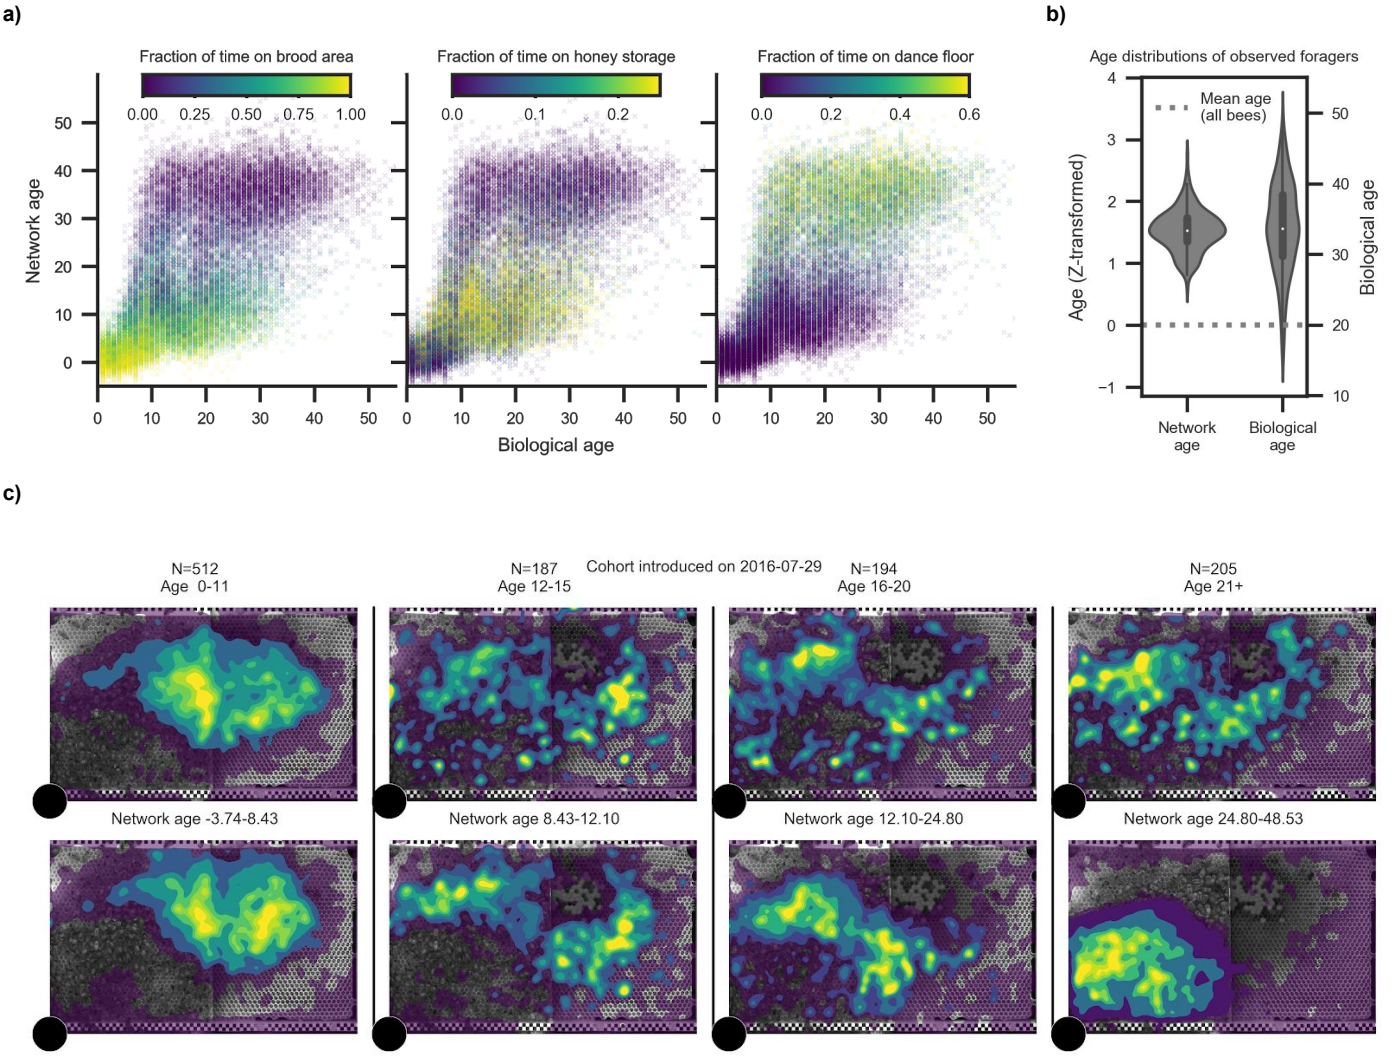

FIG 2: Network age is an accurate descriptor of task allocation. (a) The proportion of time spent on task-associated locations in relation to biological age and network age. For a given value on the y-axis (network age) colors are more consistent than for a given value on the x-axis (biological age). (b) Age distributions for known foragers visiting a feeder (N=40). Both age distributions have been z-transformed for comparison. The variance in biological age is greater than the variance in network age. Corresponding biological ages are also shown on the right y-axis (original biological age: 34.2 ± 7.9, original network age: 38.3 ± 4.6). (c) Spatial distributions of an example cohort over time (bees emerged on 2016-07-29, 64 individuals over 25 days), grouped by biological age (top row) versus network age (bottom row). Note how network age more clearly delineates groups of bees than biological age, with bees transitioning from the brood nest (center of the comb), to the surrounding area, to the dance floor (lower left area). See S008 for further examples.

Developmental changes over the life of a bee

Network age reveals differences in interaction patterns and task allocation among same-aged bees (Fig 3a). After around 6 days of biological age, the network age distribution becomes bimodal (see SI 204). Bees in the functionally old group (high network age) spend the majority of their time on the dance floor, whereas same-aged bees in the functionally young group (low network age) are found predominantly in the honey storage area (Fig 3a). Transitions from high to low network age are a rare occurrence in our colony (see S205).

We attribute the split on the population level to distinct patterns of individual development. Clustering the time series of network ages over the lives of bees identifies distinct developmental paths within same-aged cohorts (see S202 for the method). In the cohort that emerged on 2016-08-01, the first developmental cluster (blue, Fig 3b) rapidly transitions to a high network age (likely corresponding to foraging behavior) after only 11 days. The second cluster (orange) transitions at around 21 days of biological age, while bees in the third cluster (green) remain at a lower network age throughout the focal period. We see similar splits in developmental trajectories for all cohorts, although the timing of these transitions is variable (see S202 for additional cohorts). Such divergence in task allocation has been previously shown in bees; factors that accelerate a precocious transition to foraging include hormone titers (Robinson & Ratnieks 1987), genotype (Pankiw & Page 1999), physiology, especially the number of ovarioles (Amdam et al. 2004), and sucrose response threshold (Scheiner et al. 2004).

The transition from low to high network age over multiple days is characterized by a gradual shift in the spatial distribution (see example in Fig 3c), highlighting that an individual's task changes gradually. The network age of most bees is highly repeatable (median R=0.612 95% CI=[0.199, 0.982], see S201 for details), indicating task stability over multiple days. Both findings (gradual change over a few days and high repeatability) are consistent with the dynamics of the underlying physiological processes, such as vitellogenin and juvenile hormone, that influence task allocation and the transition to foraging (Amdam & Page 2010).

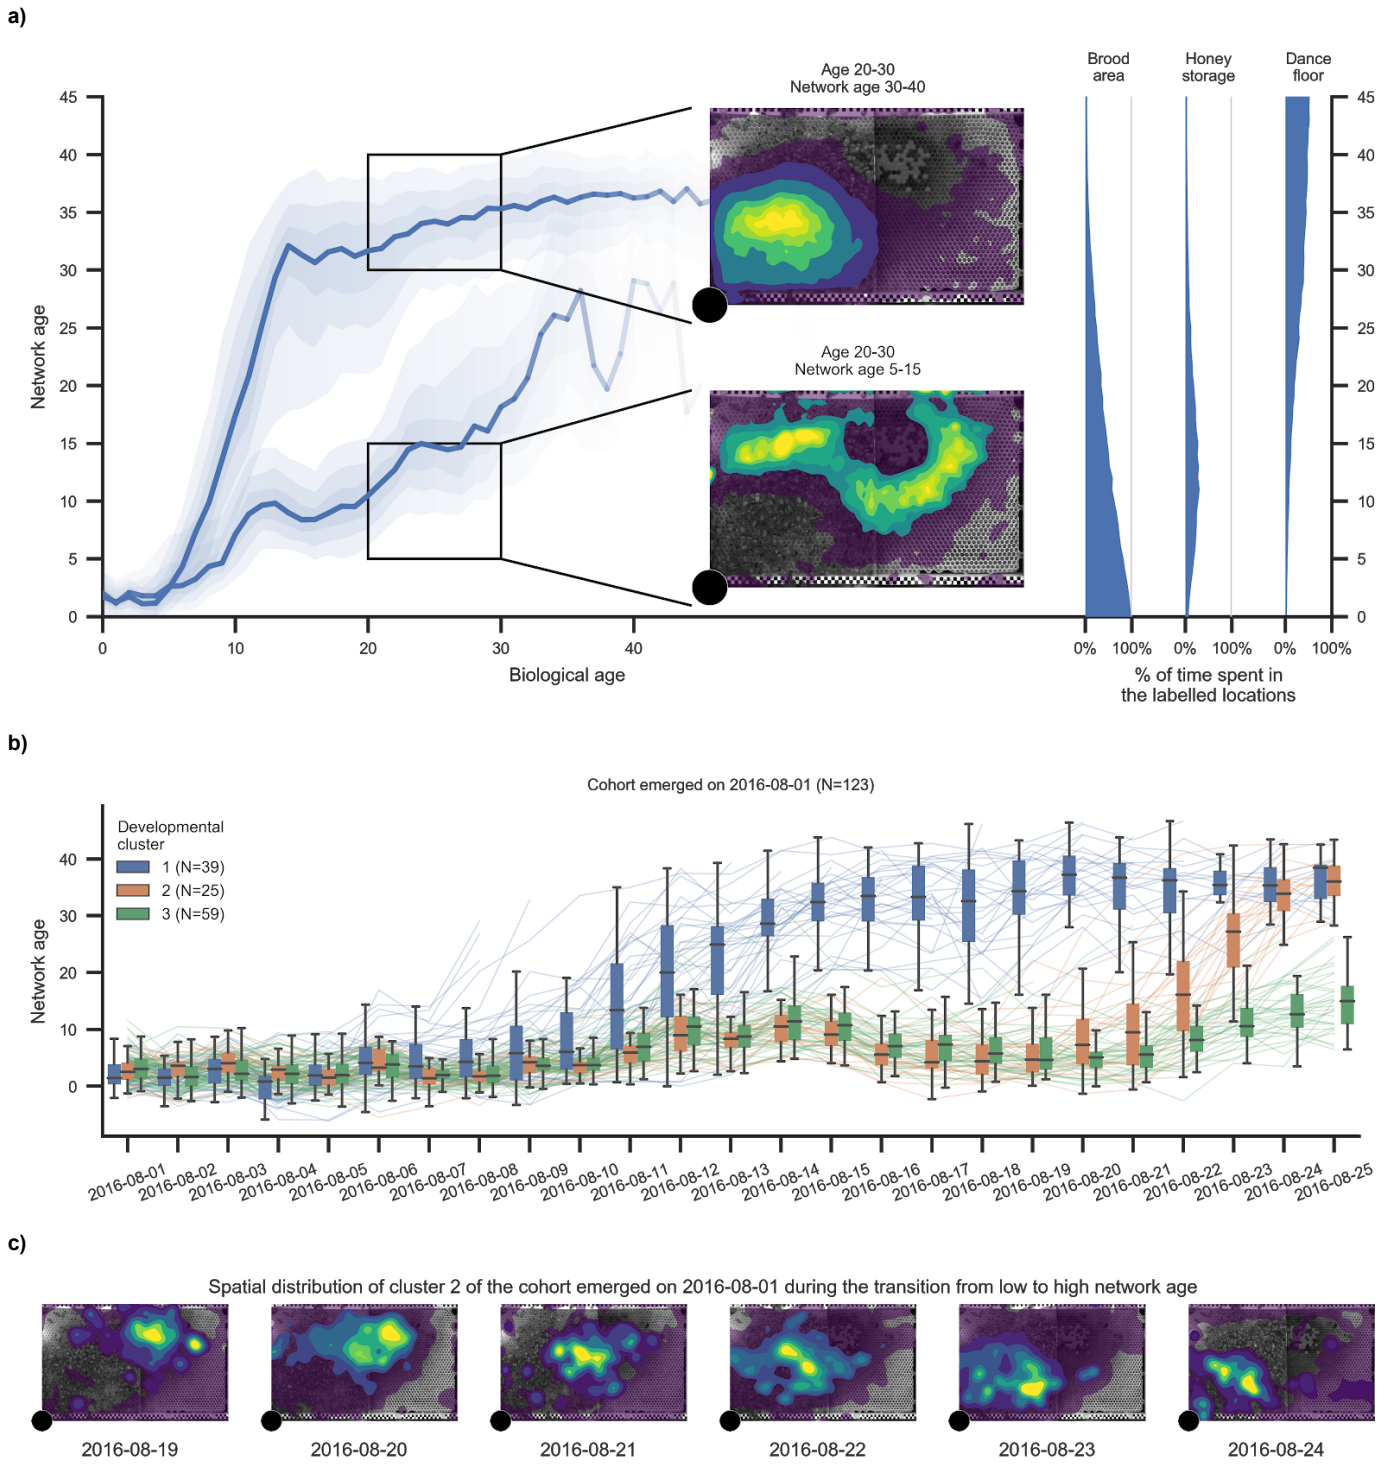

FIG 3: Bimodal distribution of network age and organization of stable "work groups" within the colony. (a) Left: The median and distribution of network ages over biological ages for all individuals that lived more than 11 days (light blue: 60% of data; see S203 for details). We observe a bimodal split in network age starting at the biological age of approximately 6 days (see S204). These two trajectories correspond to different tasks, as shown in the spatial distributions (squares indicate the bounds in network age and biological age). The upper heatmap (network age 30 to 40, biological age 20 to 30, 577 bees with 857,283 data points) corresponds to foragers on the dance floor, while the lower heatmap (network age 5 to 15, biological age 20 to 30, 381 bees with 742,622 data points) borders between the dance floor and brood nest, indicating bees likely engaged in offloading foragers and monitoring food reserves. Right: The mean fraction of time a bee with a given network age spends on our annotated regions throughout a day.

(b) The development of network age for one cohort of bees over the experimental period. Using agglomerative clustering, we grouped the bees by network age over time, which reveals different developmental trajectories. The bimodal distribution of network ages, and associated tasks, can be observed even for a single age cohort of bees. (c) Heatmaps showing the spatial distribution of bees in the developmental cluster 2 (orange) from 2016-08-19 to 2016-08-24. The smooth transition in network age (orange in line plot, b) from one mode to another corresponds to a smooth transition in spatial location (heatmaps, c). This suggests that bees change their tasks over multiple days, and that network age allows us to accurately represent these transitions.

Network age predicts an individual's behavior and future role in the colony

Network age predicts task allocation (i.e. in what part of the nest individuals will be) up to ten days into the future. Knowing the network age of a bee *today* allows a better prediction of the task performed by that individual *next week* than her biological age informs about her *current* tasks (Fig. 4C, binomial test, p << 0.001, N=55'390, 95% CI of effect size [0.055, 0.090], N=128, see S304 for details). We confirm that this is only partially due to network age being repeatable (see S304). We do note, however, that our ability to predict the future tasks of a young bee is limited, especially before cohorts split into high and low network age groups (Fig 3a). This limitation hints at a critical developmental transition point in their lives, an attractive area for future study.

We explicitly optimized network age to be a good predictor of task-associated locations. However, we find that network age predicts other behaviors better than biological age, including an individual's impending death (network age: median $R^2 = 0.165$, 95% CI [0.158, 0.172], versus biological age: median $R^2 = 0.064$, 95% CI [0.059, 0.068]; 95% CI of effect size [0.037, 0.039], N=128, likelihood-ratio χ^2 test p<<0.001, N=26,403). Biologically-young but network-old bees have a significantly higher probability of dying within a week (80.6% N=139) than do biologically-old but network-young bees (42.1% N=390; χ^2 test of independence: p<<0.001 N=529; see S305 for details). This is likely because a biologically-young bee with a high network age, i.e. a bee that starts to forage earlier in life and faces more perils imposed by the outside world, is more likely to die than a bee of the same age with a low network age. This finding is consistent with previous work showing increased mortality with precocious foraging (Rueppell et al. 2008, Perry et al. 2015).

Network age also captures movement patterns of individual bees such as daily and nightly average speed, the circadian rhythm, and the time of an individual's peak activity better than biological age (for all dependent variables likelihood-ratio χ^2 test p << 0.001, N=26'403, see Table S303A for 95% CI of effect sizes). Network age thus explains behaviors of individuals independent of their location.

While we focus on predicting tasks from network age, we can control the information we extract from the observed social networks and derive variants of network age better suited for other research questions. By replacing the *task associated location preferences* in the final step of our method with *days until death*, we extracted a descriptor that captures social interaction patterns related to mortality. This descriptor improves the prediction of the individuals' death dates by 31% compared to network age (median increase in $R^2 = 0.05$; 95% CI [0.04, 0.06] N=128, see S302 and S303A), opening up novel social network perspectives for studies such as the risk factors of

disease transmission. Similarly, we extracted descriptors optimized to predict the movement patterns introduced in the last paragraph (for all except *Time of peak activity* likelihood-ratio χ^2 test p << 0.001, N=26'403, see Table S303A for 95% CI of effect sizes). These targeted embeddings provide precise control over the type of information we extract from the social networks and extend the network age method to address other important research questions in honey bees and other complex animal societies.

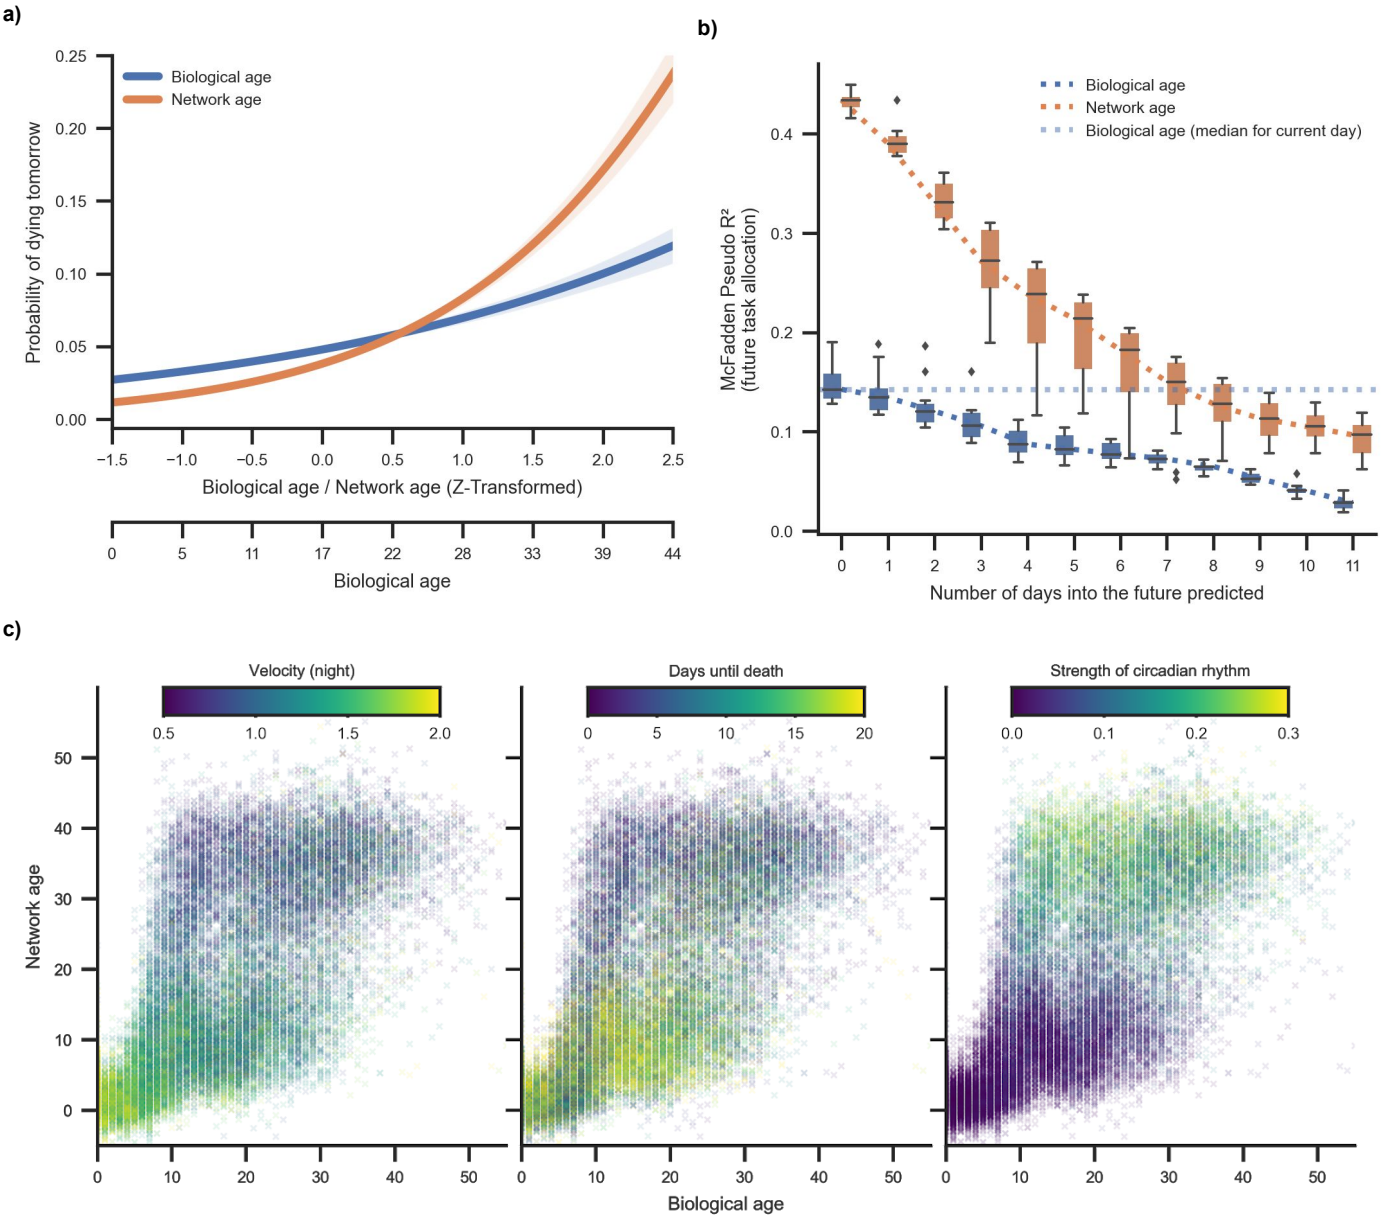

FIG 4: Network age can be used to predict other properties, such as mortality and circadian rhythms. It also predicts an individual's future task allocation. (A) An individual's mortality on the next day based on her age (x-axes show original and Z-transformed biological age and Z-transformed network age). Bees with a low network age have lower mortality than biologically young bees; bees with a high network age have higher mortality than biologically old bees. (B) Network age can be used to predict task allocation and future behaviors. Network age predicts the task of an individual 7 days into the future better than biological age predicts the individual's task the same day (blue dotted line). (C) Selected properties

mapped for network age over biological age. Note that for a given value on the y-axis (network age) colors are more consistent than for a given value on the x-axis (biological age).

Conclusion

Combining automated tracking, social networks, and spatial mapping of the nest, we provide a low-dimensional representation of the multimodal interaction network of an entire honey bee colony. We show that: (1) an individual's interaction network (network age) is a better predictor of task allocation than biological age, (2) that cohorts of individual's separate into stable functional groups in the colony, and (3) that network age predicts several aspects of an individual's life, such as mortality, movement patterns, and future task. These findings hold for several overlapping generations. Network age is derived solely from interaction patterns, not location, so bees performing similar tasks exhibit similar interaction patterns, even when they interact in different parts of the nest. Reducing interaction patterns to a single dimension inherently loses information, but our results show that the underlying structure behind how individuals interact in a colony is relatively simple and consistent over time. This perspective reveals key transition points in the lives of individuals within the colony, and how roles within the colony change over time. Obtaining network age in real-time will permit future research to analyze how complex social animal groups use and modify these interaction patterns to adapt and react to biotic and abiotic pressures.

Acknowledgments

We are indebted to numerous students, in particular Fernando Wario Vázquez, Franziska Lojewski, Andreas Jörg, Leon Sixt, Hauke Mönck, Maria Sparenberg, Sascha Witte, Alexa Schlegel, Mathis Hocke and Andreas Berg for providing hardware and software parts of the BeesBook system. We thank Peter Knoll and Randolf Menzel for providing the honey bee colony and valuable advice. Giovanni Galizia and Jake Graving gave critical feedback which helped improve the manuscript. Our work was supported by the HPC-Service of ZEDAT (Freie Universität Berlin) and the North-German Supercomputing Alliance who generously provided computing resources and storage capacity.

Author Contributions

Conceptualization: BW, DMD, AZ, TL, MLS, KST; Methodology: BW, DMD, AZ, TL; Software: BW, DMD; Resources, supervision: TL, DB, IDC; Project administration: TL; Data curation: BW, DMD, MLS; Writing: BW, DMD, TL, MLS, KST, IDC; Visualization: BW, DMD.

Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Ethics Statement

German law does not require the approval of an ethics committee for studies involving insects.

Funding

DMD received funding from the Andrea von Braun Foundation, and the Elsa-Neumann-Scholarship. MLS is a Simons Foundation Postdoctoral Fellow of the Life Sciences Research Foundation, and received funding from the Heidelberg Academy of Science and the Zukunftskolleg Mentorship Program. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824069. This work was in part funded by the Klaus Tschira Foundation. IDC and MLS acknowledge support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2117 - 422037984 and IDC acknowledges support from NSF Grant IOS-1355061 and Office of Naval Research Grants N00014-09-1-1074 and N00014-14-1-0635. KST acknowledges support of the Wissenschaftskolleg zu Berlin.

References

Abdi, Hervé. “Singular and Generalized Singular Value Decomposition.” In Encyclopedia of Measurement and Statistics, 908–12. Thousand Oaks: SAGE Publications, Inc., 2007. https://doi.org/10.4135/9781412952644.

Amdam, Gro V., Kari Norberg, M. Kim Fondrk, and Robert E. Page. “Reproductive Ground Plan May Mediate Colony-Level Selection Effects on Individual Foraging Behavior in Honey Bees.” Proceedings of the National Academy of Sciences 101, no. 31 (August 3, 2004): 11350–55. https://doi.org/10.1073/pnas.0403073101.

Amdam, Gro V., and Robert E. Page Jr. “The Developmental Genetics and Physiology of Honeybee Societies.” Animal Behaviour 79, no. 5 (2010): 973–980.

Aplin, Lucy M., Damien R. Farine, Julie Morand-Ferron, Andrew Cockburn, Alex Thornton, and Ben C. Sheldon. “Experimentally Induced Innovations Lead to Persistent Culture via Conformity in Wild Birds.” Nature 518, no. 7540 (February 26, 2015): 538–41. https://doi.org/10.1038/nature13998.

Balbuena, M. S., J. Molinas, and W. M. Farina. “Honeybee Recruitment to Scented Food Sources: Correlations between in-Hive Social Interactions and Foraging Decisions.” Behavioral Ecology and Sociobiology 66, no. 3 (March 1, 2012): 445–52. https://doi.org/10.1007/s00265-011-1290-3.

Baracchi, David, and Alessandro Cini. “A Socio-Spatial Combined Approach Confirms a Highly Compartmentalised Structure in Honeybees.” Ethology 120, no. 12 (2014): 1167–76. https://doi.org/10.1111/eth.12290.

Belkin, Mikhail, and Partha Niyogi. “Laplacian Eigenmaps for Dimensionality Reduction and Data Representation.” Neural Computation 15, no. 6 (June 2003): 1373–96. https://doi.org/10.1162/089976603321780317.

Boenisch, Franziska, Benjamin Rosemann, Benjamin Wild, David Dormagen, Fernando Wario, and Tim Landgraf. “Tracking All Members of a Honey Bee Colony Over Their Lifetime Using Learned Models of Correspondence.” Frontiers in Robotics and AI 5 (2018): 35.

Cheney, Dorothy L., Joan B. Silk, and Robert M. Seyfarth. “Network Connections, Dyadic Bonds and Fitness in Wild Female Baboons.” Royal Society Open Science 3, no. 7 (2016): 160255. https://doi.org/10.1098/rsos.160255.

Farine, Damien R., and Hal Whitehead. “Constructing, Conducting and Interpreting Animal Social Network Analysis.” Journal of Animal Ecology 84, no. 5 (2015): 1144–1163.

Flack, Jessica C., Michelle Girvan, Frans B. M. de Waal, and David C. Krakauer. “Policing Stabilizes Construction of Social Niches in Primates.” Nature 439, no. 7075 (January 2006): 426–29. https://doi.org/10.1038/nature04326.

Frisch, Karl von. “The Dance Language and Orientation of Bees.,” 1967.

Gomez-Marin, Alex, Joseph J. Paton, Adam R. Kampff, Rui M. Costa, and Zachary F. Mainen. “Big Behavioral Data: Psychology, Ethology and the Foundations of Neuroscience.” Nature Neuroscience 17, no. 11 (November 1, 2014): 1455–63.

Gordon, Deborah M. Ant Encounters: Interaction Networks and Colony Behavior. Princeton University Press, 2010. https://books.google.de/books?hl=en&lr=&id=MabwdXLZ9YMC&oi=fnd&pg=PP2&dq=Ant+encounters:+interaction+networks+and+colony+behavior&ots=PQMn0KLgWX&sig=pqjvAKVWjCYEmEfIG3EQNWHXQgY.

Goyret, J., and W. M. Farina. “Descriptive Study of Antennation during Trophallactic Unloading Contacts in Honeybees Apis Mellifera Carnica.” Insectes Sociaux 50, no. 3 (August 1, 2003): 274–76. https://doi.org/10.1007/s00040-003-0678-0.

Hasenjager, Matthew J., William Hoppitt, and Ellouise Leadbeater. “Network-Based Diffusion Analysis Reveals Context-Specific Dominance of Dance Communication in Foraging Honeybees.” Nature Communications 11, no. 1 (January 31, 2020): 1–9. https://doi.org/10.1038/s41467-020-14410-0.

Hotelling, Harold. “Relations between Two Sets of Variables.” Biometrika 28, no. 3–4 (December 1, 1936): 321–77. https://doi.org/10.1093/biomet/28.3-4.321.

Huang, Zhi-Yong, and G. E. Robinson. “Regulation of Honey Bee Division of Labor by Colony Age Demography.” Behavioral Ecology and Sociobiology 39, no. 3 (September 1, 1996): 147–58. https://doi.org/10.1007/s002650050276.

Johnson, Brian R. “Division of Labor in Honeybees: Form, Function, and Proximate Mechanisms.” Behavioral Ecology and Sociobiology 64, no. 3 (January 2010): 305–16. https://doi.org/10.1007/s00265-009-0874-7.

Knapp, Thomas R. “Canonical Correlation Analysis: A General Parametric Significance-Testing System.” Psychological Bulletin 85, no. 2 (1978): 410–16. https://doi.org/10.1037/0033-2909.85.2.410.

Krause, Jens, Richard James, Daniel W. Franks, and Darren P. Croft. Animal Social Networks. Oxford University Press, 2015.

Levene, H. “Robust Tests for Equality of Variances, p 278–292.” Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling. Stanford University Press, Palo Alto, CA, 1960.

Lusseau, David, and M. E. J. Newman. “Identifying the Role That Animals Play in Their Social Networks.” Proceedings: Biological Sciences 271 (2004): S477–81.

Luxburg, Ulrike von. “A Tutorial on Spectral Clustering.” Statistics and Computing 17, no. 4 (December 1, 2007): 395–416. https://doi.org/10.1007/s11222-007-9033-z.

McFadden, Daniel. “Conditional Logit Analysis of Qualitative Choice Behavior.,” 1973.

Mönck, Hauke Jürgen. “Large Scale Supercomputer Assisted and Live Video Encoding with Image Statistics.” Masters Thesis, Freie Universität Berlin, 2016.

Naug, Dhruba. “Structure of the Social Network and Its Influence on Transmission Dynamics in a Honeybee Colony.” Behavioral Ecology and Sociobiology 62, no. 11 (September 1, 2008): 1719–25. https://doi.org/10.1007/s00265-008-0600-x.

Nieh, James C. “A Negative Feedback Signal That Is Triggered by Peril Curbs Honey Bee Recruitment.” Current Biology 20, no. 4 (February 23, 2010): 310–15. https://doi.org/10.1016/j.cub.2009.12.060.

Otsu, Nobuyuki. “A Threshold Selection Method from Gray-Level Histograms.” IEEE Transactions on Systems, Man, and Cybernetics 9, no. 1 (January 1979): 62–66. https://doi.org/10.1109/TSMC.1979.4310076.

Pankiw, T., and R. E. Page Jr. “The Effect of Genotype, Age, Sex, and Caste on Response Thresholds to Sucrose and Foraging Behavior of Honey Bees (Apis Mellifera L.).” Journal of Comparative Physiology A 185, no. 2 (August 1, 1999): 207–13. https://doi.org/10.1007/s003590050379.

Paszke, Adam, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, et al. “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” 2019, 12.

Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, et al. “Scikit-Learn: Machine Learning in Python.” Journal of Machine Learning Research 12, no. Oct (2011): 2825–30.

Perry, Clint J., Eirik Søvik, Mary R. Myerscough, and Andrew B. Barron. “Rapid Behavioral Maturation Accelerates Failure of Stressed Honey Bee Colonies.” Proceedings of the National Academy of Sciences 112, no. 11 (March 17, 2015): 3427–32. https://doi.org/10.1073/pnas.1422089112.

Pinter-Wollman, N., E. A. Hobson, J. E. Smith, A. J. Edelman, D. Shizuka, S. de Silva, J. S. Waters, et al. “The Dynamics of Animal Social Networks: Analytical, Conceptual, and Theoretical Advances.” Behavioral Ecology 25, no. 2 (March 1, 2014): 242–55. https://doi.org/10.1093/beheco/art047.

Psorakis, Ioannis, Stephen J. Roberts, Iead Rezek, and Ben C. Sheldon. “Inferring Social Network Structure in Ecological Systems from Spatio-Temporal Data Streams.” Journal of the Royal Society, Interface 9, no. 76 (November 7, 2012): 3055–66. https://doi.org/10.1098/rsif.2012.0223.

Robinson, Gene E. “Regulation of Division of Labor in Insect Societies.” Annual Review of Entomology 37, no. 1 (1992): 637–665.

Robinson, Gene E., and Francis LW Ratnieks. “Induction of Premature Honey Bee (Hymenoptera: Apidae) Flight by Juvenile Hormone Analogs Administered Orally or Topically.” Journal of Economic Entomology 80, no. 4 (1987): 784–787.

Rosenthal, Sara Brin, Colin R. Twomey, Andrew T. Hartnett, Hai Shan Wu, and Iain D. Couzin. “Revealing the Hidden Networks of Interaction in Mobile Animal Groups Allows Prediction of Complex Behavioral Contagion.” Proceedings of the National Academy of Sciences 112, no. 15 (April 14, 2015): 4690–95. https://doi.org/10.1073/pnas.1420068112.

Rueppell, Olav, Robyn Linford, Preston Gardner, Jennifer Coleman, and Kari Fine. “Aging and Demographic Plasticity in Response to Experimental Age Structures in Honeybees (Apis Mellifera L).” Behavioral Ecology and Sociobiology 62, no. 10 (May 6, 2008): 1621. https://doi.org/10.1007/s00265-008-0591-7.

Salvatier, John, Thomas V. Wiecki, and Christopher Fonnesbeck. “Probabilistic Programming in Python Using PyMC3.” PeerJ Computer Science 2 (April 6, 2016): e55. https://doi.org/10.7717/peerj-cs.55.

Scheiner, Ricarda, Robert E. Page, and Joachim Erber. “Sucrose Responsiveness and Behavioral Plasticity in Honey Bees (Apis Mellifera).” Apidologie 35, no. 2 (2004): 133–142.

Schneider, Stanley S., and Lee A. Lewis. “The Vibration Signal, Modulatory Communicationand the Organization of Labor in Honey Bees, Apis Mellifera.” Apidologie 35, no. 2 (March 2004): 117–31. https://doi.org/10.1051/apido:2004006.

Seeley, Thomas D. “Adaptive Significance of the Age Polyethism Schedule in Honeybee Colonies.” Behavioral Ecology and Sociobiology 11, no. 4 (December 1, 1982): 287–93. https://doi.org/10.1007/BF00299306.

Seeley, Thomas D. The Wisdom of the Hive: The Social Physiology of Honey Bee Colonies. Cambridge, Mass: Harvard University Press, 1995.

Seeley, Thomas D., P. Kirk Visscher, Thomas Schlegel, Patrick M. Hogan, Nigel R. Franks, and James A. R. Marshall. “Stop Signals Provide Cross Inhibition in Collective Decision-Making by Honeybee Swarms.” Science 335, no. 6064 (January 5, 2012): 108. https://doi.org/10.1126/science.1210361.

Smith, M. L., M. M. Ostwald, and T. D. Seeley. “Honey Bee Sociometry: Tracking Honey Bee Colonies and Their Nest Contents from Colony Founding until Death.” Insectes Sociaux 63, no. 4 (November 1, 2016): 553–63. https://doi.org/10.1007/s00040-016-0499-6.

Sosna, Matthew M. G., Colin R. Twomey, Joseph Bak-Coleman, Winnie Poel, Bryan C. Daniels, Pawel Romanczuk, and Iain D. Couzin. “Individual and Collective Encoding of Risk in Animal Groups.” Proceedings of the National Academy of Sciences 116, no. 41 (October 8, 2019): 20556–61. https://doi.org/10.1073/pnas.1905585116.

Strandburg-Peshkin, Ariana, Danai Papageorgiou, Margaret C. Crofoot, and Damien R. Farine. “Inferring Influence and Leadership in Moving Animal Groups.” Philosophical Transactions of the Royal Society B: Biological Sciences 373, no. 1746 (May 19, 2018): 20170006. https://doi.org/10.1098/rstb.2017.0006.

The Astropy Collaboration, A. M. Price-Whelan, B. M. Sipőcz, H. M. Günther, P. L. Lim, S. M. Crawford, S. Conseil, et al. “The Astropy Project: Building an Inclusive, Open-Science Project and Status of the v2.0 Core Package.” The Astronomical Journal 156, no. 3 (August 23, 2018): 123. https://doi.org/10.3847/1538-3881/aabc4f.

Walt, Stéfan van der, Johannes L. Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D. Warner, Neil Yager, Emmanuelle Gouillart, and Tony Yu. “Scikit-Image: Image Processing in Python.” PeerJ 2 (June 19, 2014): e453. https://doi.org/10.7717/peerj.453.

Ward, Joe H. Jr. “Hierarchical Grouping to Optimize an Objective Function.” Journal of the American Statistical Association 58, no. 301 (1963): 236–44. https://doi.org/10.1080/01621459.1963.10500845.

Wario, Fernando, Benjamin Wild, Margaret J. Couvillon, Raúl Rojas, and Tim Landgraf. “Automatic Methods for Long-Term Tracking and the Detection and Decoding of Communication Dances in Honeybees.” Frontiers in Ecology and Evolution 3 (September 25, 2015). https://doi.org/10.3389/fevo.2015.00103.

Wario, Fernando, Benjamin Wild, Raúl Rojas, and Tim Landgraf. “Automatic Detection and Decoding of Honey Bee Waggle Dances.” PloS One 12, no. 12 (2017): e0188626.

Wild, Benjamin, Leon Sixt, and Tim Landgraf. “Automatic Localization and Decoding of Honeybee Markers Using Deep Convolutional Neural Networks,” February 13, 2018. http://arxiv.org/abs/1802.04557.

Supplementary Information

Supplementary Information - What is “network age”?

S001 Recording setup, data extraction, preprocessing

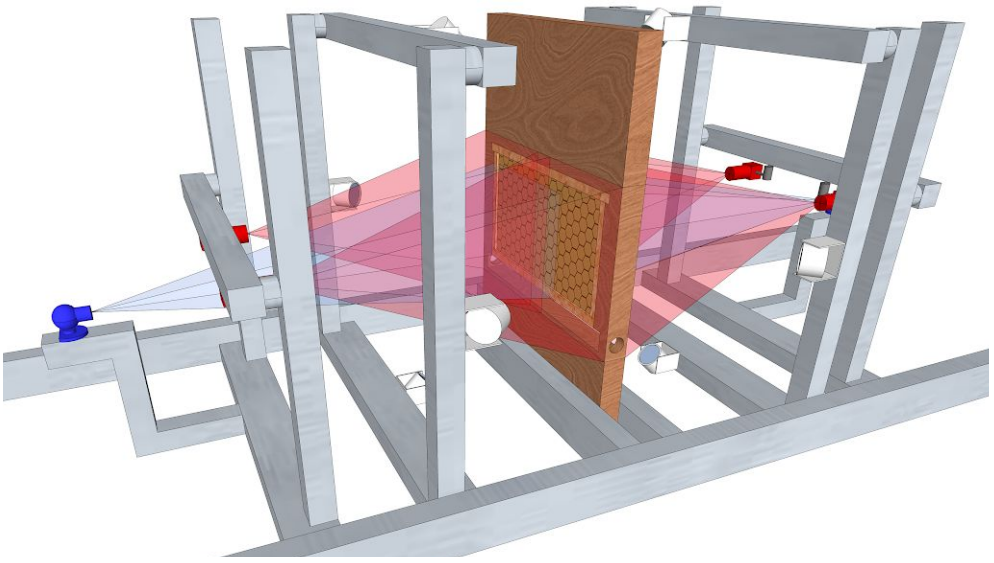

Fig S001A: Recording setup. A one-frame observation hive stands in the center of the rig, two high-res cameras (red), and one low-res camera (blue) observe each side of the hive. The low-res cameras only observe the dance floor. Four infrared flashlights were used per side (synchronized to the high-res cameras), constant red lighting illuminated the dance floor.

We set up our observation hive on 2016-06-24, with a queen and approximately 2,000 young bees (*Apis mellifera*) sourced from a local host colony. To obtain newly emerged bees, we incubated brood from the host colony, and later from the observation colony in an incubator at 34° C. Freshly emerged bees were marked every weekday. All bees were removed from the brood comb each day before marking, so the maximum age in each batch of bees was 24 hours. The number of bees marked per batch varied, but never exceeded 156 bees per marking. Marked bees were introduced to the colony through a backdoor entrance. After the initial marking period (27 days), the video recording was started 2016-07-24 and stopped 2016-09-19. Marking newborn bees continued approximately twice a week. A total of 3,166 individuals were marked. Bees had free access to the outside environment via a tube connected to the observation hive. The recording setup (scaffold, cameras, lighting, storage) was described previously (Wario et al. 2015), however, for the experiments described here, we used custom-built IR flash circuit boards triggered by an Arduino controller, synchronized with the high-res cameras (Mönck 2016). Combs were imaged at 3 Hz, alternating between sides of the observation hive, to avoid low contrast due to backlight.

A total of 46 TB of video data was recorded and continuously moved to a network storage unit at the North-German Supercomputing Alliance (HLRN). After the recording season, the data was processed to detect and decode the bee markers (Wild et al. 2018) and track these detections through time (Boenisch et al. 2018). The output data consists of timestamps, planar positions, three-dimensional rotations, IDs, and confidence scores, for the decoded IDs.

Two cameras were used for each side of the nest with both viewports overlapping partially. Six reference points were marked on each comb side such that four points were visible in each of the cameras' recordings with the center two points visible in both cameras. Reference points were identified and their image coordinates were extracted manually. The coordinates were then used to calculate the homography between the comb and the image plane. The homography was then used to rectify the tracking data which translated image coordinates to a metric reference frame, i.e. the nest surface. See Fig S001A for a photo of the setup.

Resulting tracking data was post-processed before entering the analysis. We discarded detections with low decoding confidence, i.e. detections that the machine vision pipeline could not reliably decode. Remaining implausible detections (e.g. of IDs that had not been tagged yet) were removed in an additional filter step. The distribution of the number of detections of all IDs is strongly bimodal, the larger mode representing those bees which actually are in the observation hive, and the smaller mode representing erroneous decodings of tags which are not present on the given day. We use Otsu's method (Otsu 1979) to automatically determine the threshold which best separates those two modes and filter out all potentially incorrect IDs. The tracking data, therefore, contains gaps due to falsely filtering out correct detections, but also due to occlusions, such as when bees inspect cells or depart the nest on foraging trips.

S002 Bayesian lifetime prediction - Changepoint model

The death date of an individual could ideally be computed as the first date she was not detected in the hive. Unfortunately, this does not work in practice for two reasons. First, tags are sometimes incorrectly decoded, and because of the number of detections we have for each day, this means that most IDs will be detected at least a couple of times per day. Second, some bees were not visible at all on some days, even though they are not dead yet (see Fig S002A for an example).

We, therefore, use a Bayesian changepoint model to robustly estimate the death dates of all individuals. An individual is defined to be alive on all days since she emerged and was introduced into the colony (day e) up to the change point $d = e + l$, where l is the number of days she was alive. We use a weakly informative prior $N(35,\ 50)$ for the number of alive days l. We model the probability that a bee is detected at least as often as a threshold t while she is alive and less often than t when she is not alive, using a Bernoulli distribution. We use a $Beta(5, 1)$ prior for this probability because we know that, typically, an alive bee will have many detections. For the threshold t we use an informative $Beta(25, 1)$ prior because we know that a dead bee will have very few detections, if any. Note that we normalize the detection counts to [0, 1] when fitting the

model, i.e. for each bee we divide the counts of daily detections by the maximum count of detections of that tag over the entire recording period.

We sample this model using pymc3 and the NUTS sampler (Salvatier, 2016). For each bee, we compute 2,000 tuning samples and 1,000 samples. The date of death is determined using the mean of those last 1,000 monte carlo samples.

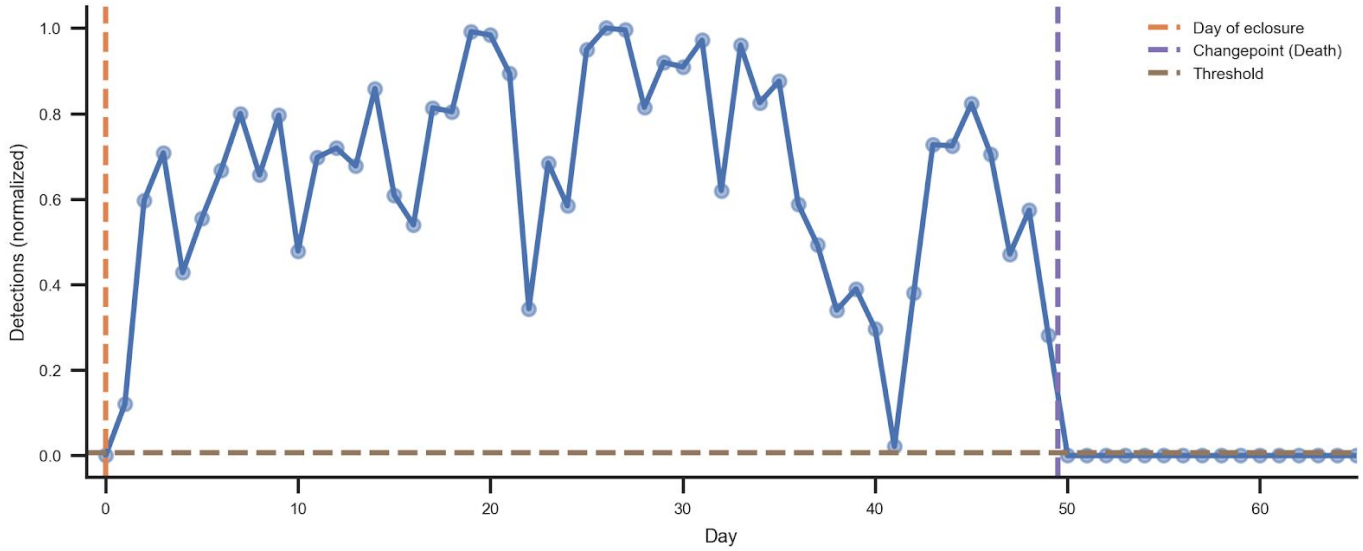

Fig S002A: The normalized detection counts of one bee over the recording period following her introduction to the observation hive. The time of her death (purple dotted line on the right) is calculated using a Bayesian changepoint model based on optimizing a detection threshold that discriminates between days that the bee was alive and days where she was probably dead.

S003 Social networks

Proximity interaction network

Two bees were defined to be in proximity if their tags were less than 2 cm from each other (~1.4 body lengths) over at least 0.9 seconds (three frames with our recording frame rate). We construct affinity networks based on the counts of these proximity interactions without taking the duration of each contact into account to reduce the effect of bees resting next to each other.

Euclidean proximity networks

Euclidean proximities were determined for each pair of bees when both bees were visible. The daily average distance d between two individuals was then transformed to two affinity matrices, the first derived by applying a gaussian similarity function $d' = \exp(-d^2/2\gamma^2)$ with $\gamma = max(D)/4$ and the second by subtracting from the maximum distance $(d' = max(D) - d)$. D is the matrix of all daily average distances on the same day.

Trophallaxis networks

We constructed an interaction network representing trophallaxis interactions (food exchange). To filter our data to detect trophallaxis events, we use a two-step approach. We first use a fast

logistic regression with low precision to discard most of the non-trophallaxis encounters. We then use a slower convolutional neural network to further refine the results with higher recall.

To train the two classification models, we manually labeled bee interactions in our dataset by observing video sequences. To increase the fraction of positive events, we queried our data for bees that are close to each other, and approximately facing each other.

This ground truth data contains 140 trophallaxis events out of the distinct 2,651 events in total. For some events, we annotated a begin and end timestamp and could, therefore, use multiple frames for the training. In total, we had 25,835 training samples, each consisting of a pair of bee IDs, a timestamp, and a label (trophallaxis / not trophallaxis).

Because the prefiltering of the training data can introduce a sampling bias we created another test set by labeling all possible interactions in 33 randomly sampled frames, containing a total of 15 trophallaxis events and 39,051 negative events (we use every pair of bees with a thorax distance of less or equal than 3cm as a possible candidate). This test set represents our data distribution without any bias.

In the classification step, we look at all pairs of bees with their thoraxes at a distance between 0.731 cm and 1.204 cm (i.e. the 99th percentile of the positive events in the training data) together in a frame. For a pair of bees (i, j) we have the locations of the thorax on the hive in millimeters (xy_i, xy_j) and their orientations (α_i, α_j). We calculate the approximate head position h_i as $xy_i + d * \left[cos(\alpha_i),\ sin(a_i)\right]$ where $d = 3.19mm$. We calculate their relative orientation as $\left[cos(\alpha_i),\ sin(a_i)\right] \cdot \left[cos(\alpha_j),\ sin(a_j)\right] T$. We then perform logistic regression on the euclidean distance of their thorax locations, the euclidean distance of their head locations, and their relative orientation.

The logistic regression was trained on the manually labeled samples, setting the threshold to get a recall of 85% at a precision of 21% (on a 20% validation test). This regression discards 62% of the true negative samples (i.e. the specificity). For the remaining data points, those that were classified as “possible trophallaxis” by the first classifier, we extract trajectories of both bees for around 5 seconds (15 frames) around the possible trophallaxis events. We then use a convolutional neural network, again trained on the manually labeled data.

Evaluated on the test set, the two combined filters yield a recall of 60% at a precision of 47%, discarding 99.97% of negative samples (i.e. the specificity).

Interaction effect networks

For each proximity interaction with a duration no longer than 60 seconds (to exclude bees resting next to each other) and with a minimum gap of at least 5 seconds since the last interaction of the same two individuals, we compute the difference in mean velocities within 30 second time windows before and after the interaction. This is done for both partners, and so we derive four networks based on the mean and cumulative changes, each split into negative and positive values. We use separate networks for the positive and negative values because this allows us to define affinity matrices which can only have positive edge weights.

Temporal aggregation and post-processing of networks

Time-aggregated networks were constructed by defining the weighted edge strength as the number of times two individuals were in proximity or engaged in trophallaxis. The networks were aggregated over 24 hours without overlap. Edges in both networks are undirected. For subsequent analyses, all networks were represented as a square adjacency matrix with each element i,j representing the affinity of bee i with bee j on this day, given by the interaction mode (e.g. for the network of trophallaxis counts, a high value represents many trophallaxis interactions between the two individuals). Each matrix is then preprocessed using a rank transform and normalized such that 0 represents the lowest affinity and 1 the highest affinity. Ties are resolved by assigning the same rank to identical affinities.

S004 Nest area mapping and task descriptor

We manually outlined the capped brood area and visible honey storage cells for every day in background images from 2016-07-30 until 2016-09-05. To obtain the open brood area, we calculated the area of the comb that would become capped within 8 days. We extracted the background images by extracting the first frame from every video we recorded over a specific day (approximately one image every 5.6 minutes), and then applying a rolling median filter with window size 10 to these images. We then calculated the modal pixel value, for every pixel, over all the median images. For each side of the comb, we stitched together the background images from the two cameras on that side.

To get the approximate location of the dance floor, we used the detected waggle runs of our waggle dance detection system (see Wario et al. 2017) that had high confidence (>=0.9) (see Fig S004A). As nearly all waggle detections happened on one side of the comb, we exclusively labeled this area as the dance floor. We then fitted an ellipse to the detections using scikit-image (van der Walt et al. 2014), which we scaled manually to not intersect with the exit area. The dance floor area was consistent throughout the experiment, so in cases where we did not have waggle dance data for a given day, we interpolated the dance floor area over the adjacent days. Finally, we used a kaiser window applied over the consecutive days (window size=5, beta=5) to smooth the annotations. We considered the region 7.5cm around the exit tube as the nest region close to the exit.

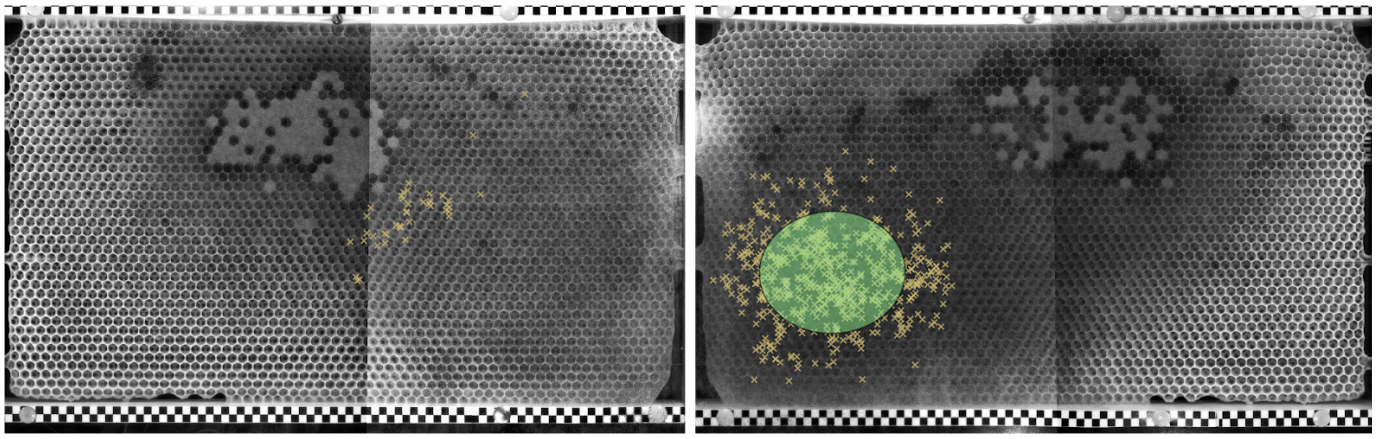

Fig S004A: Background-subtracted images of the two sides of the observation hive. The images from the two cameras per side have been stitched together. The markers show the location of the high-confidence waggle run detections for 2016-08-18. Nearly all are on one side of the comb. The green region indicates the dance floor location.

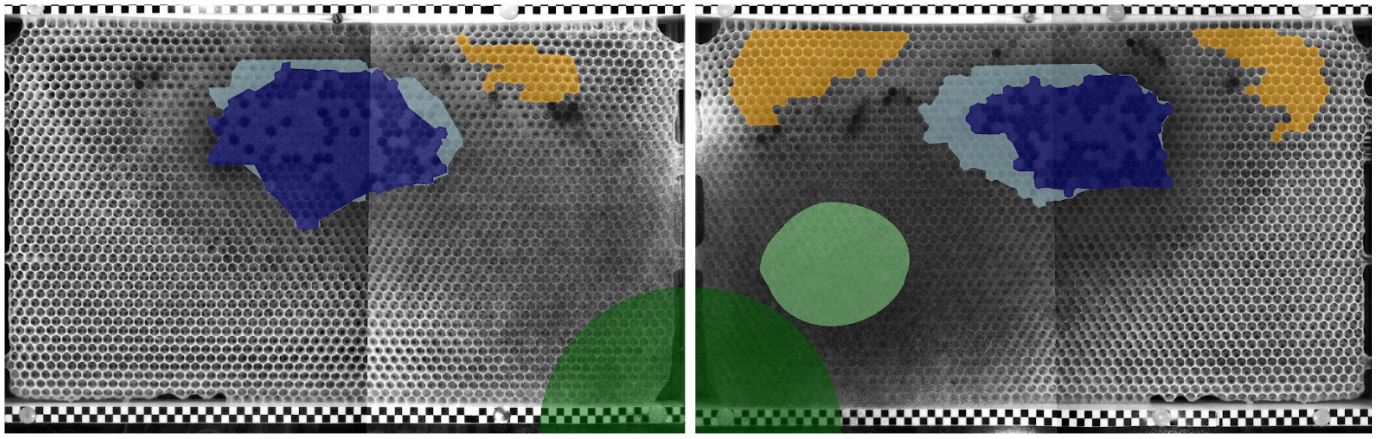

Fig S004B: Annotations for 2016-08-18, used to generate task descriptors for every bee (dark blue: capped brood, light blue: open brood, yellow: honey, light green: dance floor, dark green: exit region).

To generate a task descriptor for every bee, we fetched one high confidence detection (>0.9) per bee for every minute of a day. We then counted how many of these detections per bee fell into the annotated regions. Then we normalized these counts per bee to 1 by dividing through the sum. This descriptor, therefore, contains the fraction of time each individual spends in each of the annotated regions. Data points outside of the annotated areas are ignored for this descriptor but are used in all other parts of this work. For all evaluations, we consider the brood area region to be the sum of the annotated open and closed brood cell regions. See Fig S004B for an example of the annotations.

S005 Network age method - From networks to spectral embeddings to CCA

Network age is derived from the raw interaction matrices using spectral decomposition and canonical correlation analysis (CCA). For each day and interaction mode, the graph of interactions between all bees that were alive (see S002 for the definition of alive bees) on that day is retrieved as an adjacency matrix as described in S003.

For each preprocessed affinity matrix, spectral embeddings (Belkin & Niyogi 2003) are calculated using the python package scikit-network. We compute the first eight embedding dimensions (see S102 for an evaluation of the performance of different numbers of embeddings).

For the non-symmetric interaction effect matrices, we use bispectral decomposition (Abdi 2007) to obtain one set of embeddings each for the rows and for the columns, to represent the two directions of an interaction.

For different days the eigenvectors of the embeddings and therefore the embedding values themselves can have an inverted sign. To correct this, we flip the sign of the values if the spearman correlation between consecutive days is negative.

For every day, we now have a high-dimensional embedding per bee. We reduce the dimensionality further by applying canonical-correlation analysis (CCA). We use CCA to find a linear transformation of the network embeddings to a three-dimensional vector that maximizes the correlation to a projection of the bees' task descriptors (as introduced in S004). We use the CCA implementation in scikit-learn (Pedregosa et al. 2011). We use the first dimension of this vector as "network age" throughout this paper, but also evaluate multidimensional variants (Network age 2D, Network age 3D) in S104, S106, and S303.

For every dimension and day, we use robust scaling based on the 5th and 95th percentile of the network age distribution. A network age of zero corresponds to the 5th percentile and 40 corresponds to the 95th percentile. This stabilizes the distribution over time and also maps the values to a range comparable with the biological age of honey bees. We note that this scaling slightly improves the prediction of task allocation, but that the method also works without it. We enforce that the 5th percentile of network age corresponds to bees with a lower biological age than the 95th percentile such that biological age and network age have the same directionality.

S006 Network age method - Unsupervised variant using Principal Component Analysis

We also calculate a variation of network age that does not require the annotated location descriptors. Instead of applying CCA to the concatenated spectral embeddings (see S005), we instead use Principal Component Analysis (PCA) to reduce the dimensionality. This unsupervised network age still predicts task allocation better than biological age (see S101 and S104 for details).

S007 Network age distributions over time

While the network age distribution stays within similar bounds over time, we find that the shape of the distribution (e.g. the proportion of individuals in the lower vs higher percentiles) changes significantly over the experimental window, likely reflecting changes in task allocation (e.g. more individual engaged in brood care initially and more foragers after most brood cells are capped). See Fig S007A for distributions over time.

We find no strong evidence of correlations of these changes in distribution with external factors such as the weather, but this may be due to the limited size of the recording window and is an interesting direction for potential future research.

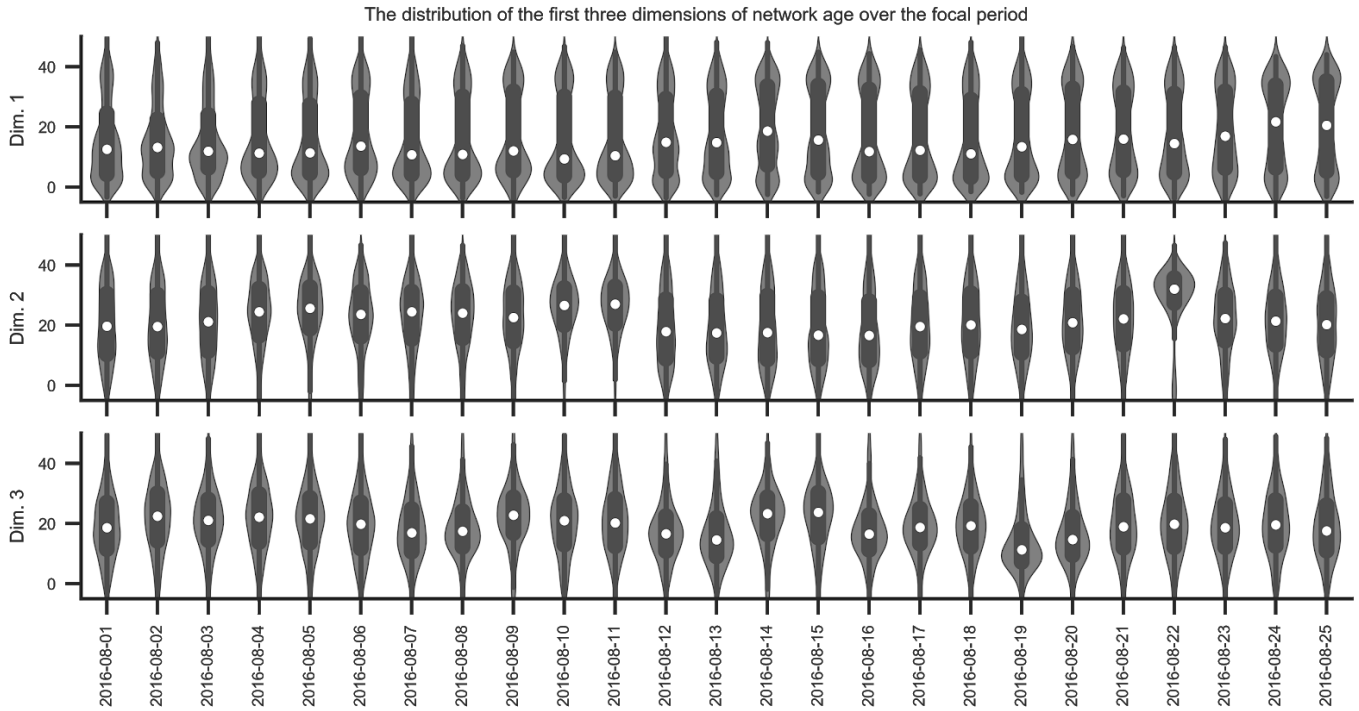

Fig S007A: Violin plots of network age distributions over the focal period. The first three dimensions of network age are depicted.

S008 Location heatmaps

To visually assess how distinctive biological age and network age are with respect to the spatial distribution of individuals, we computed location heatmaps for single-age cohorts tracked over the experimental window (25 days). For each individual within the cohort, we collected their daily biological age, network age, and their positions on the comb (as in S004). Positional data for each day (a “bee-day”) was then assigned to a heatmap as follows: For the age heatmaps, we manually defined age intervals based on the age-thresholds given by Seeley (1982), with an additional split for middle-aged bees for which we observed a high variance. We then assigned bee-days according to these thresholds based on the bees’ ages on the different days. The network age thresholds for the plots have been set in a way to keep the number of bee-days the same for the two plots in each column and can, therefore, differ for different cohorts.

For example, the cohort introduced on 2016-08-01 consisted of 123 bees. In the heatmap for age 0-11, there would be N=1394 bee-days over the focal period in which the bees were below 12 days old (some bees having died). Then, we calculated the first network age thresholds in such a way, that they would also include 1394 bee-days.

The shaded areas depict density percentiles (brightest to darkest: 99%, 97.5%, 95%, 80%, 70%, 20%).

See Fig S008A, B for heatmaps showing the distributions of different example cohorts that we introduced at the beginning of the focal period and Fig S008C for all bees.

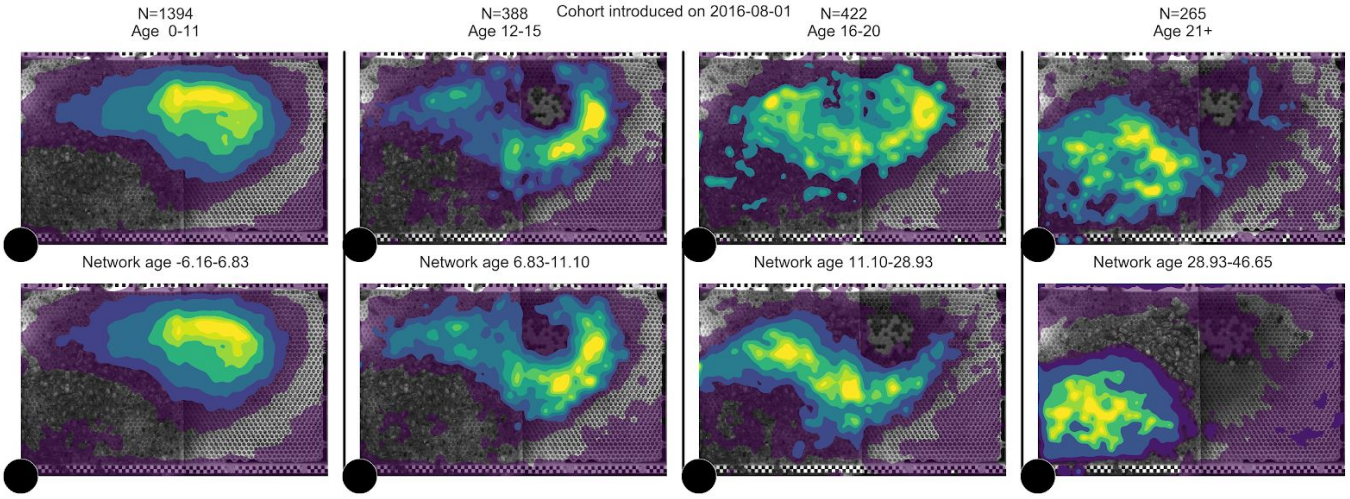

Fig S008A: Spatial heatmaps for the cohort of 2016-08-01 (N=123 bees)

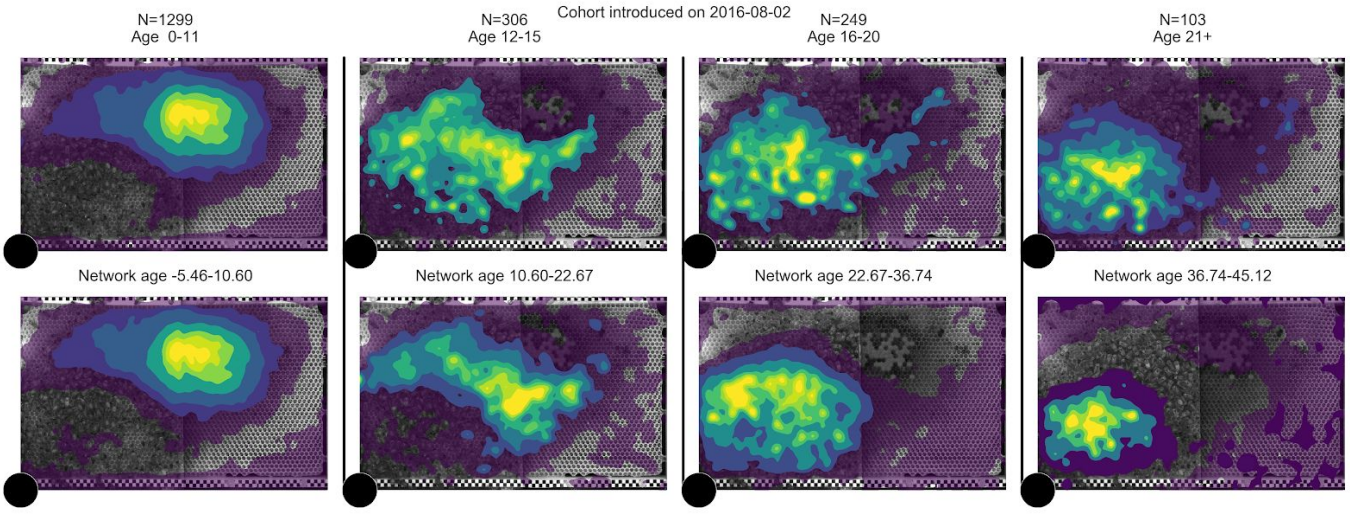

Fig S008B: Cohort from 2016-08-02 (N=119 bees)

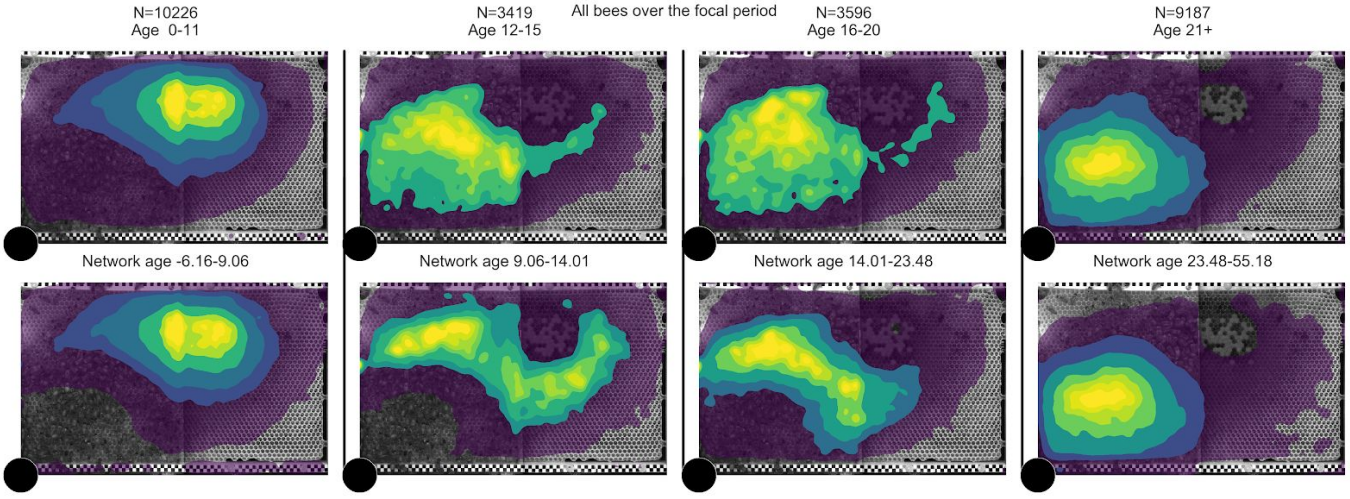

Fig S008C: All bees over the focal period (N=1912 bees)

Supplementary Information - Network age correctly identifies task allocation

S101 Variability of task prediction accuracy over time

The task allocation prediction scores are not constant over time for biological age, network age, and the unsupervised (PCA) variant of network age. Here, we show McFadden's Pseudo R^2 scores for individual days. While there is variance in the predictability using network age, we note that on all dates, the models based on network age yield a higher score than the models using biological age. On all days, the unsupervised PCA variant is more predictive than biological age with comparably small differences to the CCA variant. See Fig S101A for the distribution and timeline.

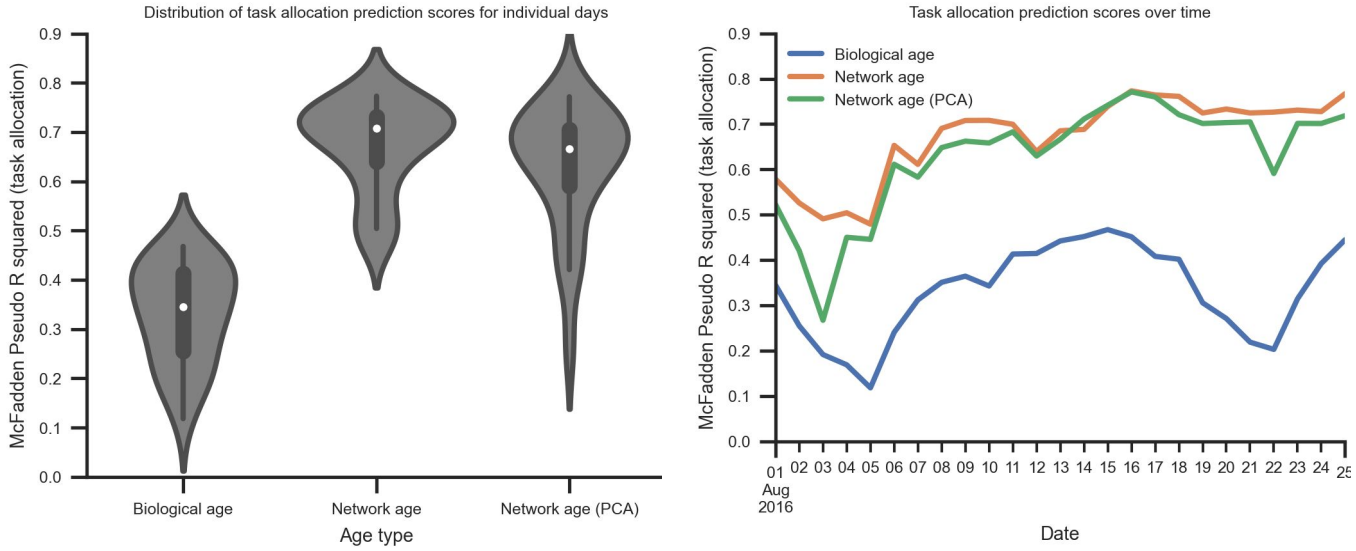

Fig S101A: (Left) The distribution of McFadden's R^2 for task allocation prediction for the different days using either network age, the unsupervised network age (PCA), or biological age as predictors. (Right) The same data points plotted over time.

S102 Effect of subsampling bees, the dimensionality of spectral embeddings, and only using proximity information on task prediction accuracy

Network age can be computed with a much smaller proportion of marked bees, hence allowing one to study much larger colonies, or reducing the effort for marking individuals. To quantify the effect of sparse sampling, we perform the following analysis. We subsample the individuals entering the analysis and compute interaction networks only for those selected individuals, i.e. only interactions between individuals of this subset are recorded. We then compute spectral factors and network age as described in S005. We find that network age is remarkably robust in this setting. With 5% of the bees tagged, the proportion in variance explained using network age

is comparable to the case of using all individuals. Even with only 1% of the individuals used in the analysis, network age still is a better predictor of task allocation than biological age. This suggests that a much larger colony could be analyzed with this method without additional effort, as long as a representative subset of the bees is tracked. See Fig S102A for the influence of different subsampling fractions on the results.

Likewise, we also find that the number of spectral factors per interaction mode matrix does not have a strong effect on the proportion of variance explained by the network age extracted from those factors. Using more factors is computationally more expensive and produces statistically better embeddings. However, network age computed using 32 factors explains task allocation only marginally better compared to network age computed using 4 factors (see Fig S102A).

We also evaluate variants of network age that can be derived using only the spatial information of detected individuals. To that end, we exactly follow the procedure outlined in S005, but discard all interaction types except for either spatial proximities ("Euclidean proximity" variant, see "Euclidean proximity networks" in S003) or interaction events derived from spatial proximity ("Proximity events" variant, see "Proximity interaction network" in S003). We find that network age derived from all interaction types outperforms these proximity based variants in terms of task allocation, but the variant based on proximity events is only marginally worse (see S104) and could be used in future studies if only proximity data is available.

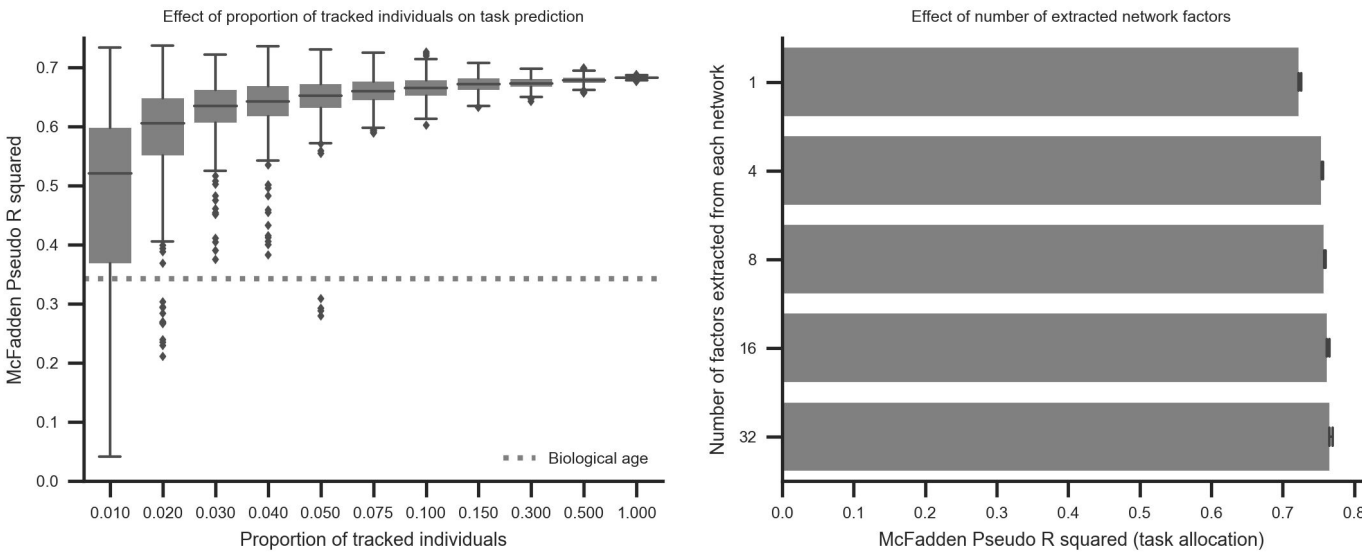

Fig S102A: (Left) McFadden's R^2 for different fractions of subsampled individuals. While having all bees individually marked and tracked yields the highest median score, tracking only 5% of the individuals yields comparable results when drawing a representative subset. (Right) The influence of the number of dimensions per interaction network extracted using Spectral embedding before applying the CCA on the task allocation prediction scores (using network age 3D). While more dimensions enable the CCA to extract more relevant information, the overall improvement is small. The confidence intervals are based on 128 bootstrap samples per dimensionality.

S103 Spatial separation of bees with similar network age

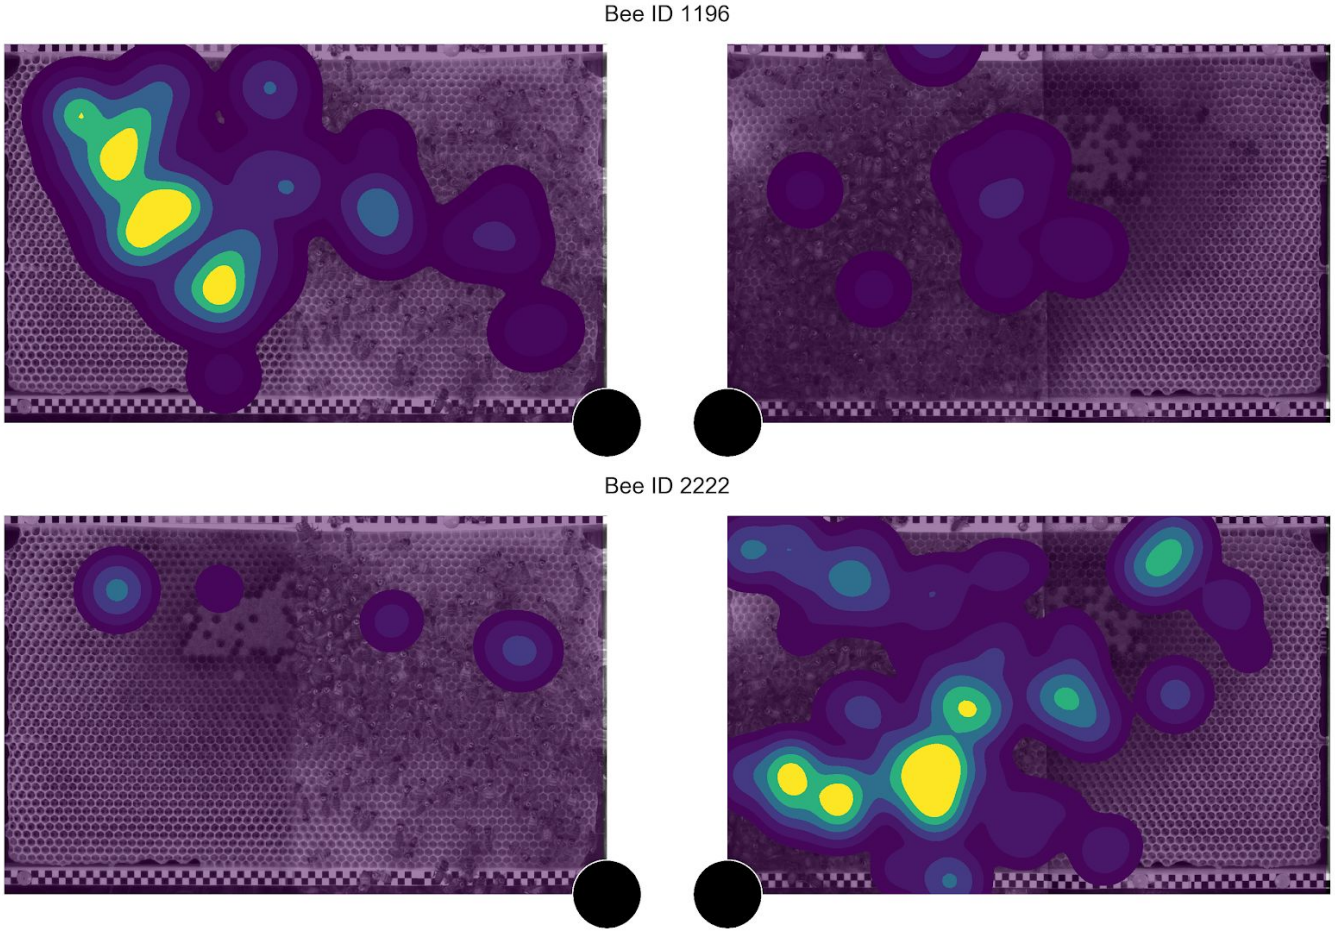

Fig S103A: Two bees on 2016-08-05 with similar network ages despite occupying different sides of the comb. The bee with the ID 1196 is 22 days old and has a network age of 15.09. Her heatmap consists of 1189 samples, each corresponding to one minute on the focal day. The bee with the ID 2222 is 8 days old and has a network age of 14.34. Her heatmap consists of 1284 samples, each corresponding to one minute on the focal day. The black circle marks the location of the entrance.

While the network age of bees is correlated with their locations, we do find examples of bees that have approximately the same network age but occupy different locations on the comb (see Fig S103A for an example). We also see that, while two bees having a similar network age is correlated with an increase in proximity interactions, there are many bees with close network ages that do not encounter each other inside the hive (see Fig S103B).

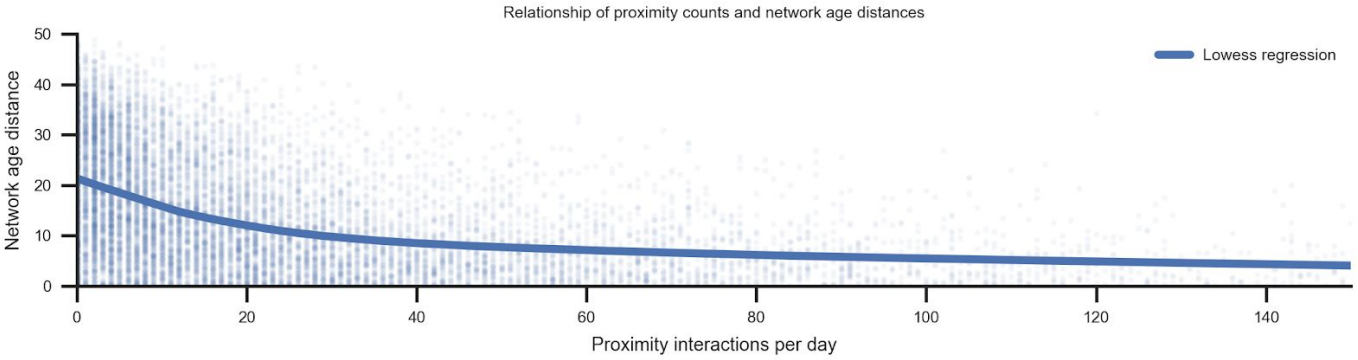

Fig S103B: The relationship between distance in network ages and number of proximity interactions for pairs of bees. Close network ages generally mean more interactions. However, there are many individuals with a close network age and very few interactions.

S104 Task prediction models and bootstrapping

To evaluate how well biological age and the different variants of network age represent an individual's task allocation, we use these measures as features to predict the proportion of time individuals spend in the brood area, dance floor, honey storage and near the exit (see S004 for details on the nest area mapping). We evaluate the areas individually and in combined models. We evaluate different complexities of models (linear vs. nonlinear) and different independent variables (e.g. network age and biological age).

To test different complexities of the relationships, we evaluate both a generalized linear model (GLM, the default model) and a small neural network consisting of two fully connected layers (listed as 'nonlinear' in the results table) for each of the combinations of independent and dependent variables. The hidden layer of the neural network has a dimensionality of 8 and uses tangens hyperbolicus as its nonlinearity.

To evaluate the performance of the models for each area separately we select a sigmoid as the link function of the GLM and the activation function of the neural network's last layer. We then optimize and calculate the likelihood of the data assuming a binomial distribution.

We also evaluate both models to simultaneously predict all four values of an individual's task allocation distribution. To this end, we choose a softmax function as the link function of the GLM and the neural network's final activation function. We then optimize and calculate the likelihood of the data assuming a multinomial distribution.

For all the combinations of independent and dependent variables, we repeat the described procedure for 128 bootstrap samples. For each model, we retrieve the final likelihood of the data L_1. We use PyTorch (Paszke et al. 2019) and the L-BFGS optimizer to obtain maximum likelihood estimates of the models. We also always fit a null-model only consisting of the intercepts and retrieve its likelihood L_0. For each model and bootstrap iteration, we calculate McFadden's pseudo R² (McFadden 1973) as $R^2_{McF} = 1 - (\log(L_1)/\log(L_0))$. We then calculate the median and 95% confidence intervals from these bootstrap samples. See Table S104A for an

overview of the results for all evaluated models. We test the significance of these results separately with the tests described in S106.

Dependent variables	Brood area	Dance floor	Honey storage	Near exit	Combined
Age	0.461 [0.452, 0.472]	0.276 [0.268, 0.286]	0.002 [0.001, 0.003]	0.324 [0.313, 0.334]	0.342 [0.335, 0.349]
Age (nonlinear)	0.508 [0.499, 0.518]	0.352 [0.343, 0.360]	0.091 [0.085, 0.096]	0.414 [0.403, 0.422]	0.384 [0.377, 0.390]
Network age (1% of bees tagged)	0.584 [0.070, 0.794]	0.505 [0.079, 0.685]	0.030 [0.000, 0.207]	0.491 [0.116, 0.734]	0.516 [0.135, 0.705]
Network age (Euclidean proximities)	0.511 [0.503, 0.519]	0.372 [0.362, 0.380]	0.166 [0.157, 0.176]	0.532 [0.523, 0.542]	0.518 [0.512, 0.524]
Network age (1% of bees tagged, nonlinear)	0.622 [0.081, 0.811]	0.579 [0.161, 0.728]	0.162 [0.032, 0.340]	0.554 [0.166, 0.772]	0.546 [0.193, 0.718]
Network age (Euclidean proximities, nonlinear)	0.603 [0.596, 0.610]	0.550 [0.542, 0.558]	0.197 [0.176, 0.209]	0.563 [0.553, 0.574]	0.553 [0.547, 0.560]
Network age (PCA)	0.796 [0.791, 0.800]	0.529 [0.519, 0.540]	0.002 [0.002, 0.004]	0.577 [0.569, 0.585]	0.646 [0.641, 0.650]
Network age (5% of bees tagged)	0.764 [0.656, 0.825]	0.607 [0.522, 0.683]	0.032 [0.005, 0.080]	0.601 [0.521, 0.686]	0.650 [0.578, 0.705]
Network age (PCA, nonlinear)	0.800 [0.795, 0.804]	0.607 [0.599, 0.616]	0.273 [0.260, 0.282]	0.657 [0.650, 0.663]	0.660 [0.656, 0.665]
Network age (5% of bees tagged, nonlinear)	0.771 [0.629, 0.831]	0.657 [0.581, 0.727]	0.240 [0.151, 0.322]	0.668 [0.571, 0.732]	0.668 [0.595, 0.721]
Network age (Proximity events)	0.793 [0.788, 0.797]	0.644 [0.636, 0.652]	0.031 [0.028, 0.034]	0.596 [0.586, 0.603]	0.674 [0.669, 0.678]
Network age	0.801 [0.797, 0.806]	0.634 [0.626, 0.642]	0.034 [0.030, 0.038]	0.611 [0.602, 0.620]	0.682 [0.678, 0.687]
Age + Network age	0.806 [0.802, 0.810]	0.634 [0.626, 0.642]	0.101 [0.093, 0.109]	0.614 [0.605, 0.622]	0.691 [0.688, 0.694]
Network age (Proximity events, nonlinear)	0.808 [0.802, 0.811]	0.679 [0.672, 0.687]	0.316 [0.306, 0.325]	0.680 [0.673, 0.686]	0.699 [0.694, 0.703]
Network age (nonlinear)	0.816 [0.812, 0.821]	0.680 [0.674, 0.687]	0.325 [0.317, 0.334]	0.694 [0.686, 0.700]	0.708 [0.704, 0.712]
Age + Network age (nonlinear)	0.821 [0.818, 0.825]	0.687 [0.682, 0.694]	0.356 [0.346, 0.366]	0.696 [0.689, 0.703]	0.716 [0.712, 0.718]
Network age (2D)	0.843 [0.839, 0.846]	0.647 [0.639, 0.653]	0.457 [0.446, 0.465]	0.638 [0.630, 0.645]	0.734 [0.730, 0.737]
Network age (2D, nonlinear)	0.850 [0.847, 0.853]	0.693 [0.686, 0.698]	0.492 [0.482, 0.501]	0.701 [0.696, 0.707]	0.744 [0.740, 0.747]
Network age (3D)	0.843 [0.840, 0.847]	0.705 [0.697, 0.712]	0.459 [0.448, 0.467]	0.703 [0.696, 0.709]	0.759 [0.755, 0.762]
Network age (3D, nonlinear)	0.857 [0.853, 0.860]	0.747 [0.740, 0.752]	0.502 [0.493, 0.509]	0.762 [0.756, 0.767]	0.769 [0.765, 0.772]

Table S104A: McFadden's R^2 median scores and bootstrapped 95% CI for different combinations of independent variables in the rows and dependent variables in the columns. We evaluated the following independent variables: different dimensionalities of network age ("Network age", "Network age (2D)", "Network age (3D)", see S005), network age derived without supervision ("PCA", see S006), with a subsample of the bees ("1% of bees tagged" and "5% of bees tagged", see S102), and based only on spatial proximity ("Euclidean proximities" and "Proximity events"). The likelihood ratio test was evaluated for the combined models (e.g. "Age + Network age", see S106).

S105 Forager groups experiment and data recording

Between 2016-07-28 and 2016-08-22, foragers were trained to a feeder offering unscented sucrose solution by gradually moving it from the colony (52.457032, 13.296635) to a sequence of locations F1 to F4 (see Table S105A). For days over which the feeder was moved, high sugar concentration was used and iteratively changed to control the number of new foragers. Once the final locations were reached, the feeder offered the highest concentration for 1-2 hours per day. After a minimum of three days, training to the next location in the list was resumed. We photographed all bees landing at the location and manually transcribed the identity and time of arrival. A list of foragers visiting the feeder is given in Table S105B.

Feeder	GPS coordinate	Distance from hive
F1	52.455726, 13.294224	217 m
F2	52.454866, 13.293092	338 m
F3	52.454547, 13.291737	430 m
F4	52.453682, 13.292841	451 m

Table S105A: GPS coordinates and distances from the hive for the feeders used in the forager group experiment.

Date	Location	Forager bees observed at feeder
2016-08-01	Hive-F1	135, 199, 217, 220, 228, 233, 253, 319, 392, 539, 727, 818, 845, 860, 1110
2016-08-02	F1	135, 217, 233, 253, 648, 845, 860
2016-08-03	F1	135, 228, 233, 253, 392, 648, 845, 860
2016-08-04	F1	220, 228, 233, 392, 648, 710, 885
2016-08-05	F1	135, 220, 228, 233, 392, 644, 648, 710, 769, 885, 931
2016-08-08	F1-F2	199, 319, 337, 392, 555, 644, 710, 769, 885, 1593
2016-08-09	F1-F2	199, 319, 337, 392, 555, 644, 648, 710, 769, 785, 885, 912, 1122, 1362, 1518, 1593, 2031
2016-08-10	F1-F2	199, 319, 337, 392, 648, 769, 885, 1122, 1362, 1593, 2031
2016-08-11	F2	199, 319, 337, 392, 644, 769, 785, 885, 1122, 1362, 1518, 1593, 2031
2016-08-12	F2	199, 319, 337, 392, 769, 785, 885, 1122, 1362, 1518, 1593, 2031
2016-08-14	F2	199, 319, 337, 392, 555, 644, 710, 769, 785, 885, 1232, 1362, 1518, 1593, 2031
2016-08-16	F2	199, 319, 337, 710, 769, 885, 912, 1122, 1362, 1518, 1593, 2031
2016-08-17	F2	199, 319, 337, 644, 710, 769, 885, 912, 1122, 1518, 1593, 2031
2016-08-18	F2	199, 319, 337, 644, 769, 885, 912, 1122, 1518, 2031
2016-08-19	F2	319, 710, 885, 1593, 2031
2016-08-22	F2-F3	319, 644, 710, 885, 1593, 2031
2016-08-23	F3	1180, 1197, 1232, 1471, 1593, 1662, 1714, 1799, 2031, 2106, 2984
2016-08-24	F3	1180, 1197, 1232, 1471, 1593, 1714, 1793, 1799, 2031, 2106, 2984
2016-08-25	F4	1180, 1197, 1714, 1793, 1799, 2031, 2106

Table S105B: Tag IDs and locations observed forager bees during the forager group experiment.

Fig S105A: Individually marked bees at a feeding site.

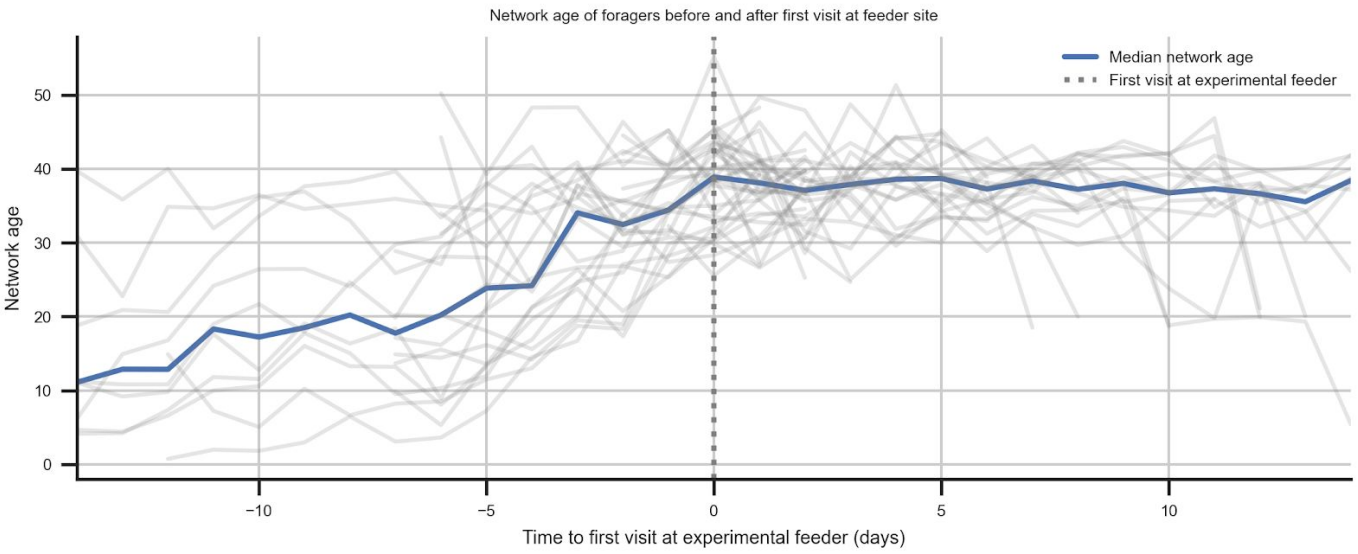

Fig S105B: For every visiting bee, their first visit to a feeding station was recorded (age 12-40 days, N=40). The plot shows the network age for these bees over the previous and next days centered on their first visit.

S106 Statistical comparison of models

We use bootstrapped confidence intervals of the effect strength to investigate whether a model based on one feature (e.g. network age) explains the dependent variables (e.g. task allocation distributions) significantly better than the same model based on a different feature (e.g. biological age). Additionally, we use a likelihood ratio χ^2 test to answer whether one feature (e.g. network age) provides additional information over biological age in a combined model.

Bootstrapped confidence intervals of the effect strength

We draw 128 bootstrap samples of the combined daily bee data. For each sample, we calculate either the McFadden's pseudo R² in the case of the task allocation models (see S104, S304) or the R² in the case of the other measures (see S303, S305) for both a model based on biological age and the independent variable we want to compare with (e.g. network age). For each of these paired samples, we calculate the difference in scores of the two models. From these 128 differences, we calculate a two-sided 95% confidence interval of the effect strength. If the null hypothesis (difference in scores is zero or less) is not contained in the confidence interval, we can reject the null hypothesis at an alpha level of 2.5%.

Likelihood ratio test

As the likelihood ratio test requires a nested model for testing, we compare a model based solely on biological age with a model based on a combination of biological age and the independent variable we want to compare with (e.g. network age).
We fit each model to the data and calculate the likelihoods of the data under the fitted models (L_1 for the combined model and L_0 for the model based on biological age). The likelihood ratio is given by $LR = -2 * ln(L_0/L_1)$. If the null hypothesis that the models are equal were true, LR would

approximately follow a chi-squared distribution with k degrees of freedom (with k=4 in the case of the task allocation model from S104 and k=1 in case of the general regression model for S303, S305). We use the cumulative density function of the chi-squared distribution to calculate the p-value.

Supplementary Information - Developmental changes over the life of a bee

S201 Repeatability

We calculate the repeatability, R, of the network age consisting of repeated measurements over several days of an individual I as $R(I) = Var_p / (Var_i + Var_p)$ with Var_i being the variance of the network age of an individual measured over the available days and Var_p being the variance of the mean network ages of a control group. The control group consists of all bees inside the same age span as I on all days on which the network age values for I were collected. A repeatability close to 1 means that the individual variance is low compared to the population variance. A repeatability close to 0 means that the individual variance outweighs the population variance.

S202 Network age transition clustering

In order to cluster the transitions of different individuals in a cohort, we first collect the network ages for every individual in a feature vector where each entry corresponds to the individual's network age for one day. We then do a linear inter- and extrapolation for missing values (e.g. due to absence or the individual dying). For the cohort of bees, we calculate the euclidean distance between each individual's feature vector. Then we perform a hierarchical clustering using Ward's method (Ward 1963) using the Python library scikit-learn (Pedregosa et al. 2011) and extract the first three clusters. See Fig 202A for an example of the clustering. See Fig 202B for the network age development of different cohorts and Fig 202C for all bees.

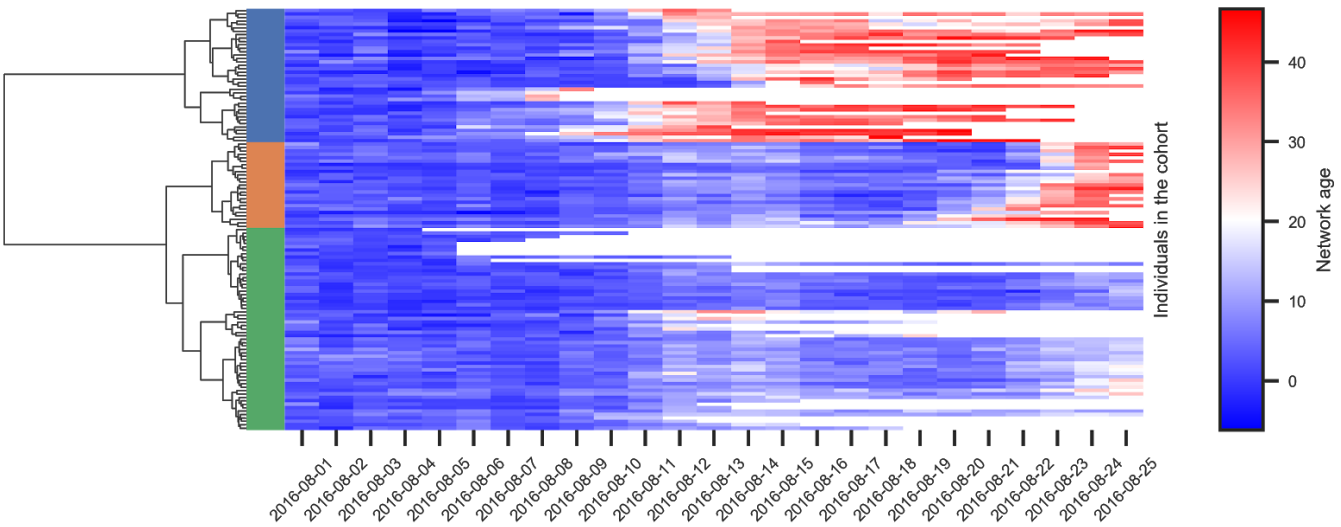

Fig S202A: Agglomerative clustering of the network age over time of one cohort of bees (emerged on 2016-08-01). Missing values (e.g. due to the death of the bee) are white.

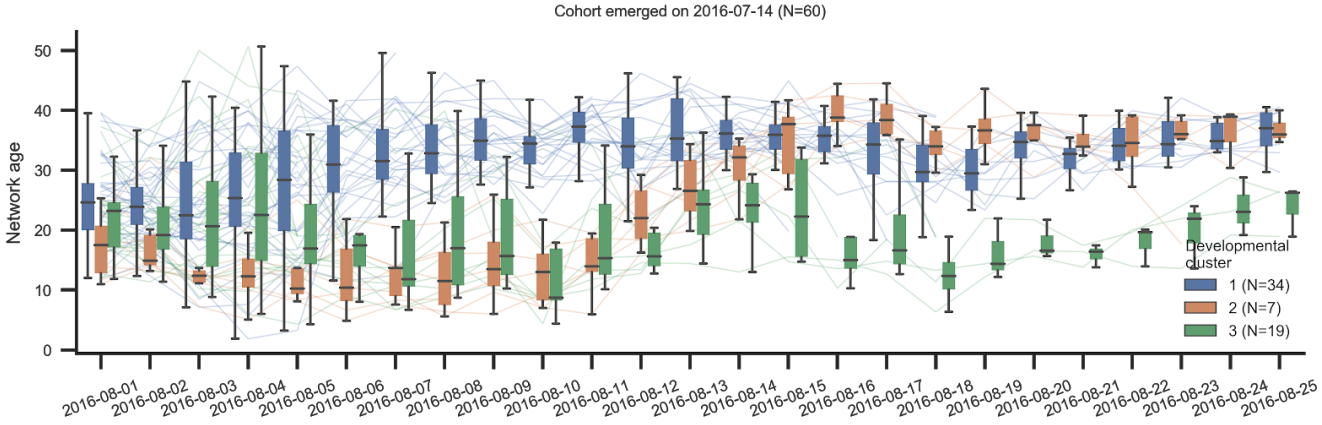

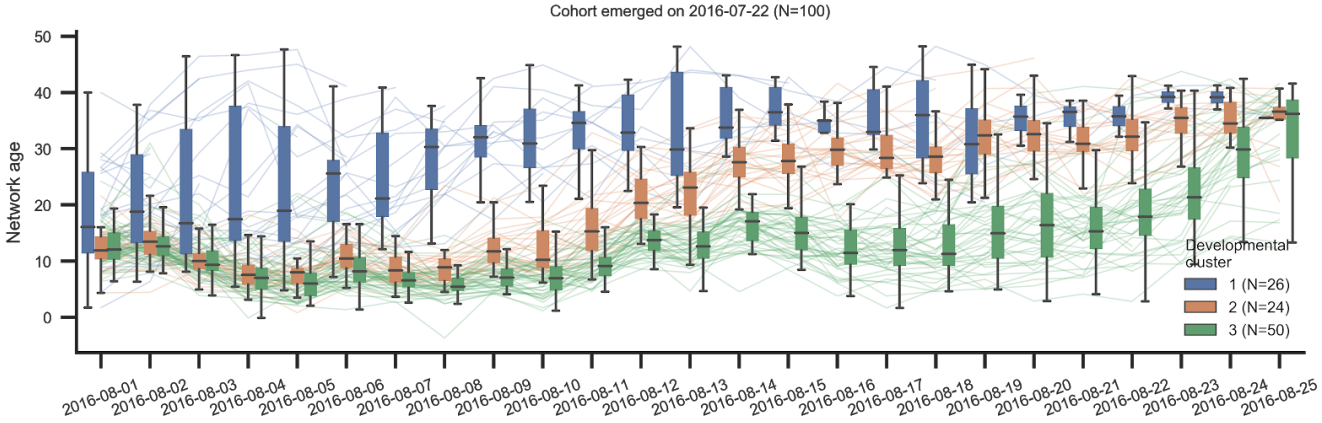

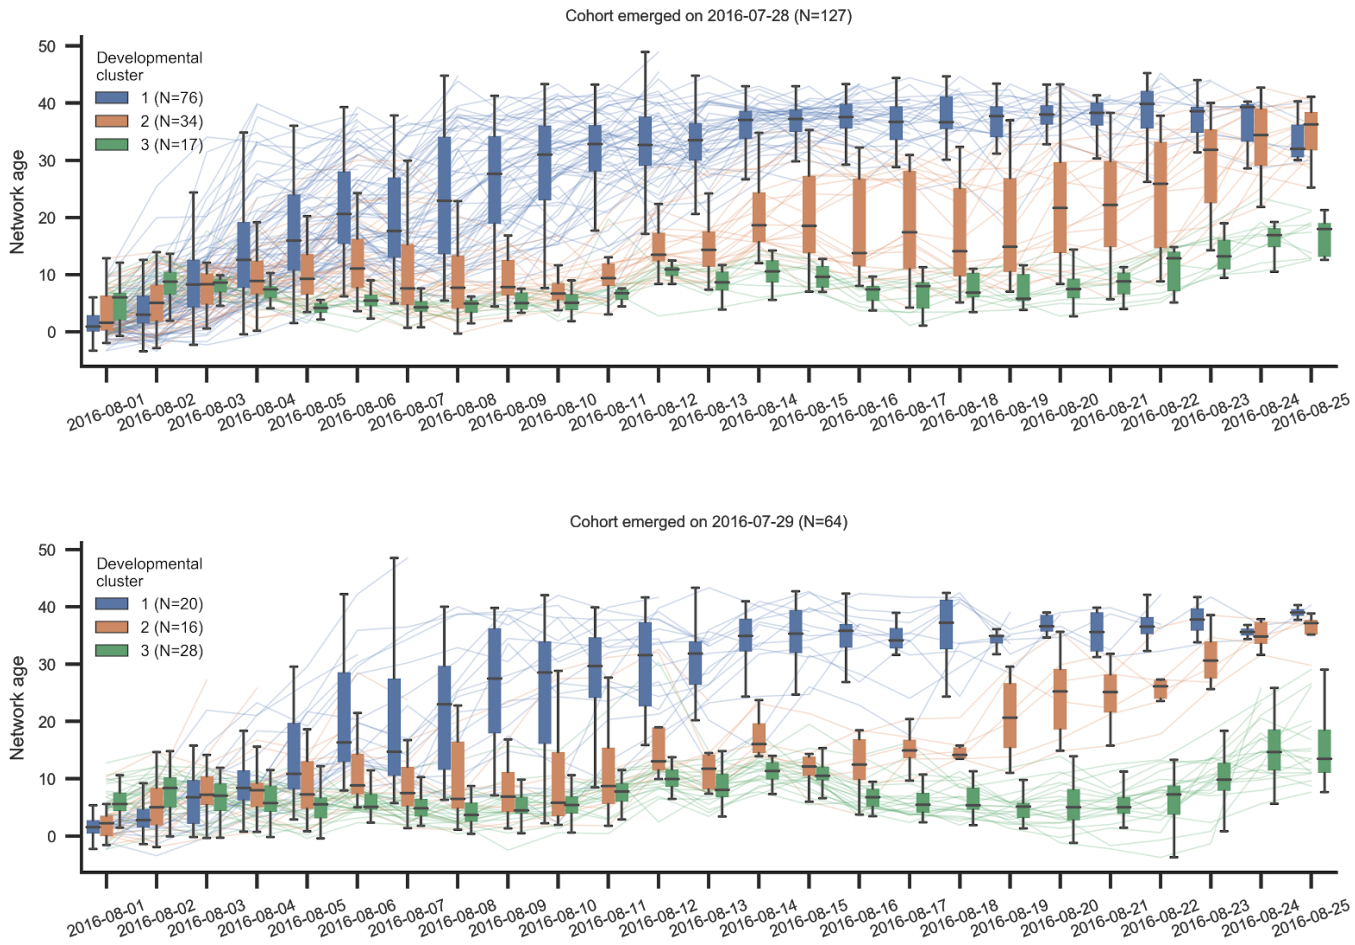

Fig S202B: The development of the network age of different cohorts (see individual titles) over the focal period. The individuals are grouped according to the agglomerative clustering.

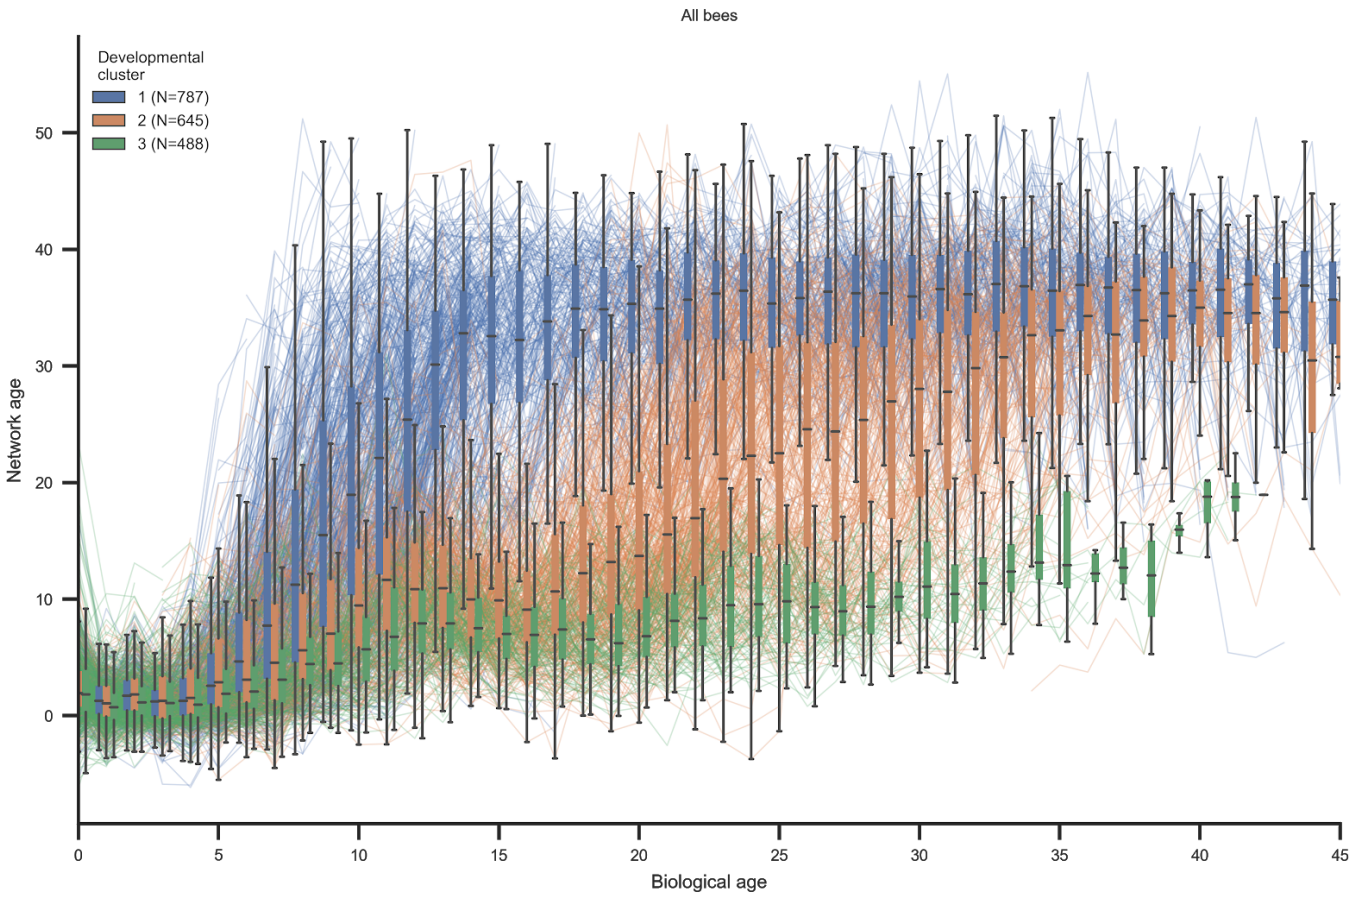

Fig S202C: The network age development of all individuals over their age. To calculate this, we followed a procedure analogous to clustering the cohorts but used network age per age (instead of per date) as the bees' feature vectors. Note that there was only a subset of 25 days of data available for each individual as some individuals were already in the hive as the focal period began and some were introduced later.

S203 Threshold for bimodal network age distribution

For the bees depicted in Fig 3a, we calculate a threshold T=23.07 on the network age using Otsu's method (Otsu 1979). The upper line contains all bees that fall above this threshold (N=832), the lower contains all bees below that threshold (N=563). The shaded areas depict 20%, 40%, 60% data intervals.

S204 Quantifying when bees first split into distinct network age modes

We used KMeans to cluster the network age distribution of every day into two distinct clusters corresponding to the two modes. Then we check for every bee that we observed at least once as a young bee below the age of 6 (N=1079) at which age she first gets assigned to the higher cluster (mean=12.33, 95% CI = [6, 25.7], median=11, N=572). We ignore bees that are never assigned to the upper cluster (e.g. because we do not observe them for a long enough timespan). See Fig 204A for the distribution of biological ages.

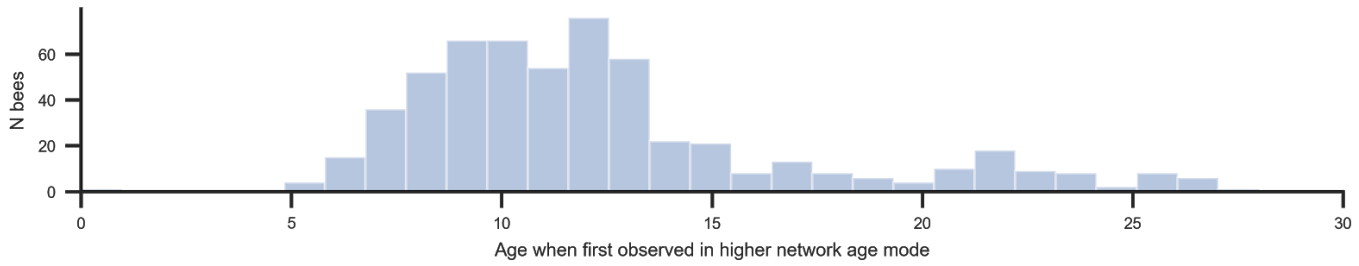

Fig 204A: The distribution of the biological age at which a bee was first assigned to the higher network age cluster. Starting at around age 5-7, the network age distribution becomes bimodal.

S205 Change of network age mode over a week

We analyze how many bees that are assigned to either of the network age modes on day X are assigned to the same cluster on day X+7 (clusters calculated as in S204). Most of the bees in the upper cluster are also in the upper cluster one week later (mean=93% 95% CI [77%, 100%], N=3698 bee days).

Bees in the lower cluster tend to stay there less (mean=56% 95% CI [0%, 96%], N=10927 bee days) depending on their age (see Fig 205A).

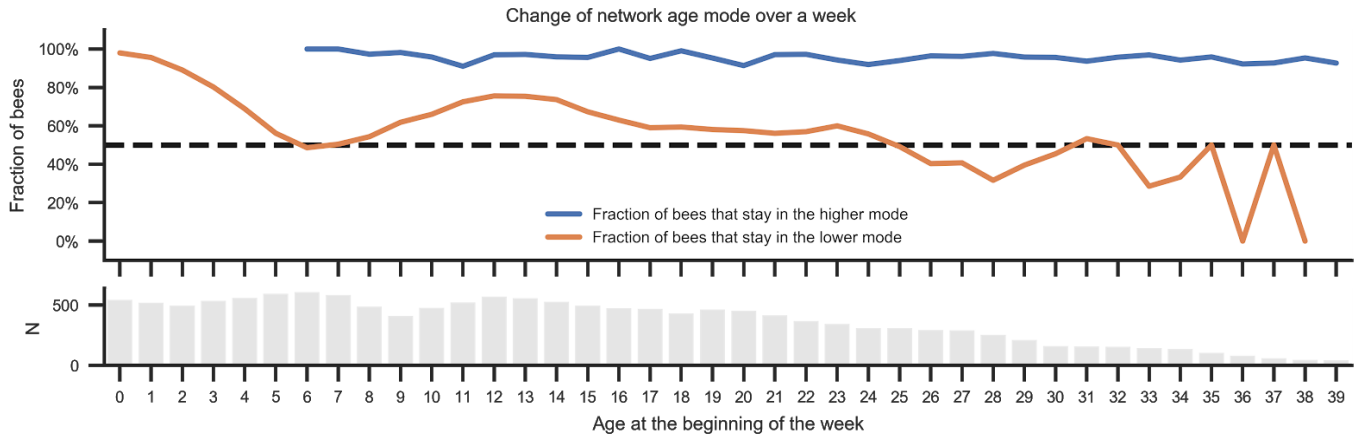

Fig 205A: (Top) We considered every individual and day (first day) for which we had data one week later (seventh day). We sorted them by their biological age on the first day (x-axis). Bees that were assigned to the higher network age cluster on the first day make up the blue line, bees from the lower cluster are depicted by the orange line. The y-axis value of each line is the fraction of bees that are still assigned to the same cluster seven days later. A high proportion of bees that are assigned to the higher cluster stay there over the course of one week (blue line). There are more bees that transition from the lower cluster to the higher one, regardless of age (orange line). (Bottom) The number of data points (orange and blue combined) for each age.

Supplementary Information - Network age predicts an individual's behavior and future role in the colony

S301 Definition of circadian rhythmicity

The motion velocity of a bee was determined by dividing the euclidean distance between two consecutive detections by the time passed (a multiple of ⅓ seconds). Any duplicate IDs were discarded. The velocity was median filtered with a kernel size of 3 to remove outliers. See Fig S301A for an example of the velocity of a specific bee over multiple days.

Lomb-Scargle periodograms were computed for all individuals remaining in the dataset at any time in the period from 2016-07-20 to 2016-09-18. For each day and individual, the Lomb-Scargle periodogram was calculated on the motion velocities over an interval of three days, i.e. including the preceding and following day. The circadian activity was confirmed as strong peaks at a period of 1 day. Lomb-scargle periodograms were computed using the Astropy package (The Astropy Collaboration et al. 2018). See Fig S301 for an example of a bee's velocity and the resulting Lomb-scargle periodogram.

In the following analyses, we reduced computational load by fitting a single sine wave of fixed frequency = 1/d (least-squares fit). For each fit, we extract the power as

$$P(f) = 1 - \frac{SSE_{sine}}{SSE_{constant}}$$

SSE_{sine} : the residuals (sum of squared errors) of the sine fit

$SSE_{constant}$: the residuals (sum of squared errors) of a constant model assuming the mean of the data

The power, hence, reflects how much of the velocity variation can be explained by the sinusoidal oscillation, or circadian rhythm.

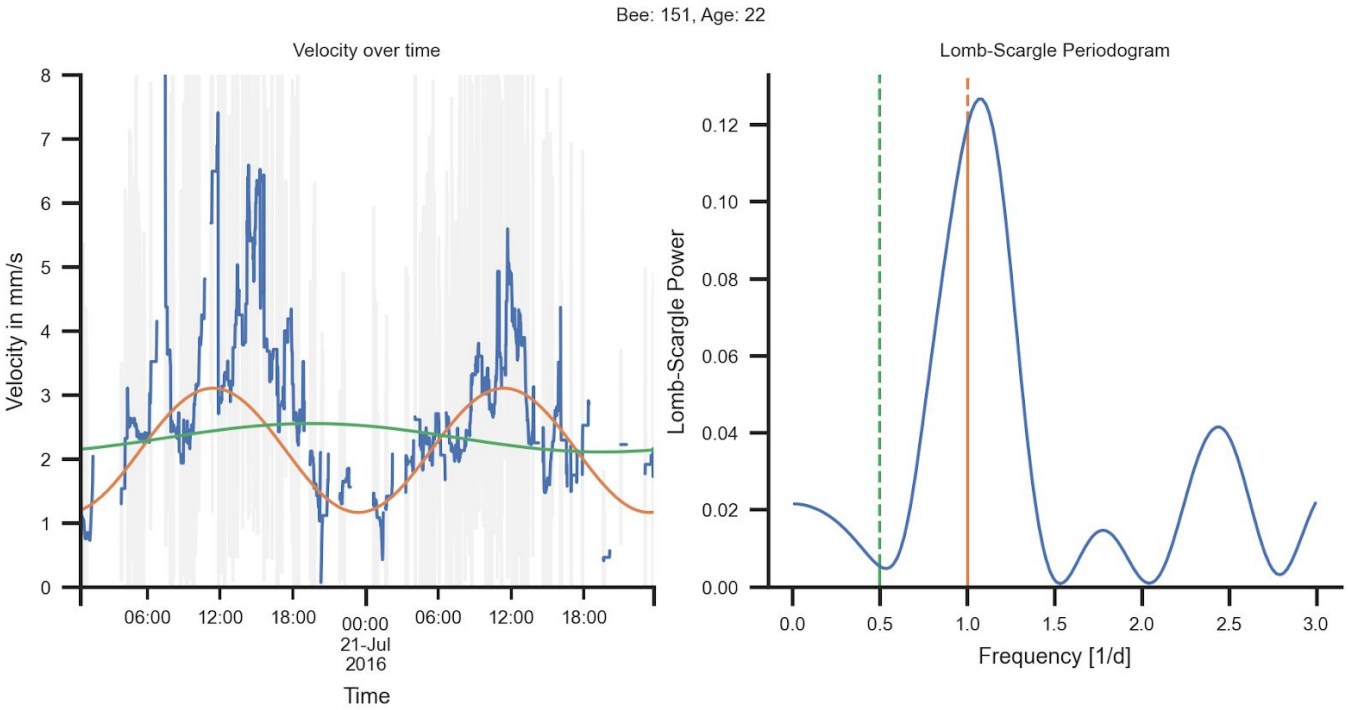

Fig S301A: (Left) The velocity as raw data (grey line) and smoothed with a median filter (blue line) for one bee over two days. The orange sine wave has a period of one day and was fitted to the velocity data via least-squares fitting. The green line is a similar fit with a period of two days. (Right) The corresponding Lomb-Scargle periodogram with the powers at the one-day frequency (orange) and two-day frequencies (green) highlighted.

S302 Targeted embedding using CCA

We show that we can extract targeted embeddings from the spectral factors of the interaction networks that are better in predicting other properties of the individuals. To compute those targeted embeddings, we exactly follow the methodology outlined in S005, but for each property (days until death, time of peak activity, circadian rhythm, day- and nighttime velocities) we extract a one-dimensional embedding from the spectral factors that maximizes the correlation with this property using canonical-correlation analysis.

S303 Prediction of other behavior-related measures

We follow the same methodology as described in S104 to evaluate how well the various age measures explain various properties of the individuals with a slight adjustment: We choose the identity function as the link function of the generalized linear model and as the activation function of the neural network's final layer. We model the residual distribution as a normal distribution with constant variance.

We define the 'days until death' as the number of days left in an individual's life on a given day (for a description of the automatically determined death dates see S002). The time of peak activity and the rhythmicity of daily movement were calculated as described in S301.

For the daytime and nighttime velocities we use the same data as for the circadian rhythmicity (see S301). For the daytime velocities we use the mean of all collected velocities between 09:00-18:00 UTC of the three-day rolling window; for nighttime velocities, 21:00-06:00 UTC.

See Table S303A for an overview of the scores of different models and targets.

(a) R^2 scores for dependent and independent variables and models

Dependent variables	Days until death	Time of peak activity	Strength of circadian rhythm	Velocity (daytime)	Velocity (night)
Age	0.064 [0.059, 0.068]	0.072 [0.065, 0.078]	0.295 [0.285, 0.302]	0.124 [0.118, 0.132]	0.215 [0.204, 0.226]
Age (nonlinear)	0.074 [0.069, 0.078]	0.078 [0.067, 0.087]	0.310 [0.300, 0.317]	0.142 [0.134, 0.150]	0.346 [0.335, 0.358]
Network age	0.165 [0.158, 0.172]	0.103 [0.097, 0.110]	0.612 [0.605, 0.619]	0.240 [0.232, 0.246]	0.322 [0.312, 0.334]
Age + Network age	0.168 [0.159, 0.175]	0.106 [0.100, 0.114]	0.612 [0.604, 0.621]	0.241 [0.233, 0.249]	0.332 [0.322, 0.343]
Network age (nonlinear)	0.197 [0.190, 0.204]	0.108 [0.102, 0.114]	0.630 [0.622, 0.636]	0.273 [0.262, 0.282]	0.372 [0.361, 0.386]
Network age (2D)	0.230 [0.222, 0.237]	0.110 [0.102, 0.117]	0.621 [0.613, 0.628]	0.259 [0.251, 0.268]	0.369 [0.355, 0.379]
Network age (3D)	0.231 [0.224, 0.240]	0.109 [0.104, 0.116]	0.628 [0.623, 0.634]	0.296 [0.289, 0.305]	0.371 [0.362, 0.383]
Targeted embedding (1D)	0.216 [0.209, 0.224]	0.103 [0.097, 0.110]	0.623 [0.615, 0.629]	0.305 [0.296, 0.315]	0.406 [0.395, 0.416]
Targeted embedding (1D, nonlinear)	0.229 [0.222, 0.238]	0.108 [0.100, 0.114]	0.635 [0.626, 0.642]	0.309 [0.299, 0.318]	0.411 [0.399, 0.421]
Network age (2D, nonlinear)	0.267 [0.258, 0.274]	0.117 [0.110, 0.124]	0.632 [0.625, 0.639]	0.289 [0.280, 0.299]	0.398 [0.387, 0.409]
Age + Network age (nonlinear)	0.227 [0.218, 0.235]	0.126 [0.116, 0.134]	0.646 [0.639, 0.653]	0.298 [0.290, 0.308]	0.436 [0.425, 0.450]
Network age (3D, nonlinear)	0.272 [0.264, 0.282]	0.122 [0.116, 0.131]	0.644 [0.639, 0.651]	0.330 [0.321, 0.339]	0.404 [0.394, 0.417]

(b) Improvement in R^2 scores (network age vs biological age)

Dependent variable	Effect size [95% CI]
Days until death	0.100 [0.093, 0.109]
Time of peak activity	0.031 [0.023, 0.040]
Strength of circadian rhythm	0.318 [0.307, 0.327]
Velocity (daytime)	0.115 [0.106, 0.124]
Velocity (night)	0.107 [0.089, 0.124]

(c) Improvement in R^2 (network age vs targeted embeddings)

Dependent variable	Effect size [95% CI]
Days until death	0.050 [0.042, 0.061]
Time of peak activity	0.000 [-0.009, 0.011]
Strength of circadian rhythm	0.009 [0.000, 0.019]
Velocity (daytime)	0.065 [0.052, 0.077]
Velocity (night)	0.083 [0.072, 0.095]

Table S303A: a) R^2 values for different models when predicting different behavior-related metrics. Biological age consistently performs worse than predictions based on network age. Our method can also be applied without spatial information by optimizing the CCA for these metrics, yielding embeddings that are strongly predictive of them while still only using information available from the social networks (see 'Targeted embedding (1D)'). (Bottom) 95% bootstrap confidence intervals of the effect sizes (b) R^2 values for network age and biological age, c) R^2 values for targeted embeddings and network age, N=128).

S304 Future predictability

We evaluate how well we can predict future task allocation using network age and biological age. To ensure that no information leak can occur, we only used supervised information from the past and test it on future data. To do this, we first calculate the spectral factors for the entire dataset as described in S005. The factors are computed for each day separately, hence no information can leak from the past to the future. We then determine the mapping from spectral factors to network age using CCA, but only on a fixed range of days prior to the validation window. We fix the number of days in the train set to 12 days. Similar to the linear mapping given by CCA, we also determine the parameters of the regression model only on the train dataset from the same fixed time window. We train separate models for all viable ranges of dates and for prediction from one to 11 days into the future. The linear mapping given by the CCA and the predictive models are then applied to the spectral factors on the held-out validation set from a time after the training data set (see Fig S304C for an overview about the data handling). For this analysis, we want to evaluate how well we can predict task allocation into the future. We estimate the effect size in

R^2_{McF} by calculating the 95% CI for the different time windows (see Fig S304C, median improvement in $R^2_{McF} = 0.080$, 95% CI [0.055, 0.090], N=12).

Because the two compared models are not nested, the likelihood ratio test does not apply here. We perform a paired binomial test using the null hypothesis that the improvement in mean squared error in task allocation prediction is zero or less. We find that we can predict the task allocation of an individual seven days into the future with a lower mean squared error (paired binomial test, p << 0.001, N=55'390) using network age.

This is mostly caused by older bees. For young bees, there is a low amount of variance in task allocation and network age is about as good in describing task allocation for young bees as biological age. Furthermore, we find that we can't reliably predict future task allocation for young bees, suggesting that either the social networks in this study are not predictive for this task or that the future task allocation of young bees is driven by other factors. See Fig S304A for an overview of the results.

The predictive power of network age cannot be fully explained by the bimodal distribution of network age and organization of stable "work groups" within the colony. We perform an additional analysis to reject this repeatability null hypothesis by comparing how well the task allocation prediction using networks works when simply shifting the prediction for the current day into the future. We find that a model fitted to predict future task allocation outperforms this null model considerably (See Fig S304B).

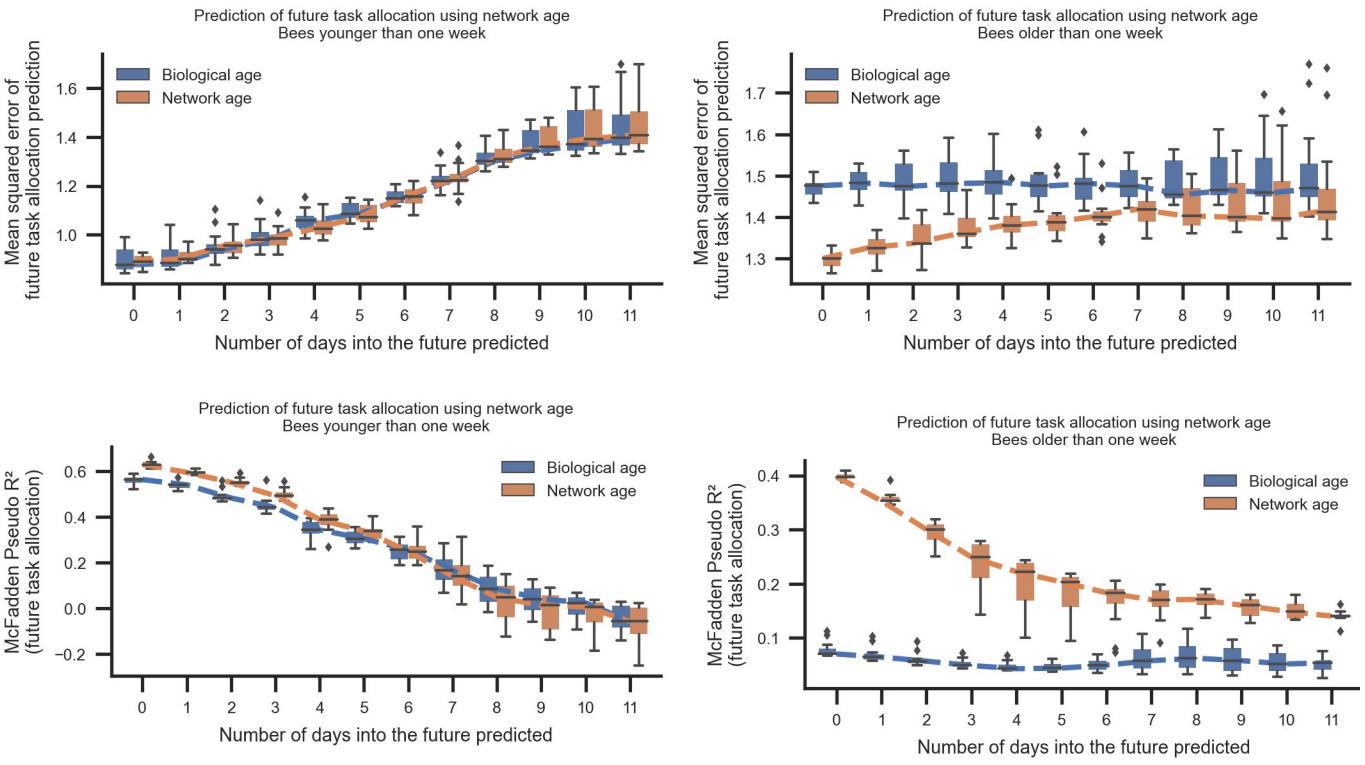

Fig S304A: Mean squared error (top row) and McFadden's R² (bottom row) for a model that predicts task allocation up to 11 days into the future. We distinguish between young bees (left column) and old bees (right column). For the old bees (> 7 days of age), network age consistently performs better than age.

When predicting task allocation 11 days into the future for old bees, network age performs better than age performs for the current date. For young bees, we find neither biological age nor network age to be predictive.

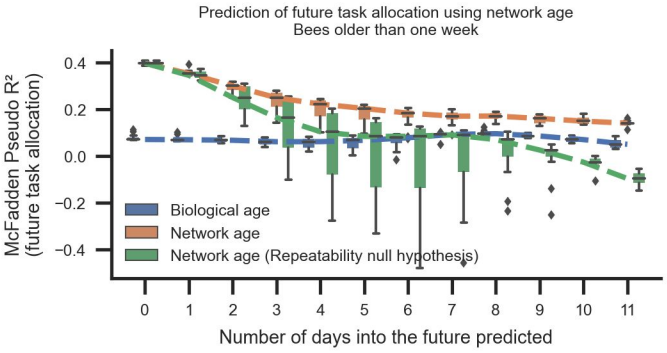

Fig S304B: To check that the predictiveness of network age for old bees is not simply due to the bimodal distribution of the network age and its high repeatability, we compare a predictive model based on network age (orange) with the null-hypothesis that the network age of day X is sufficient to describe the coming days without explicitly modeling the changes.

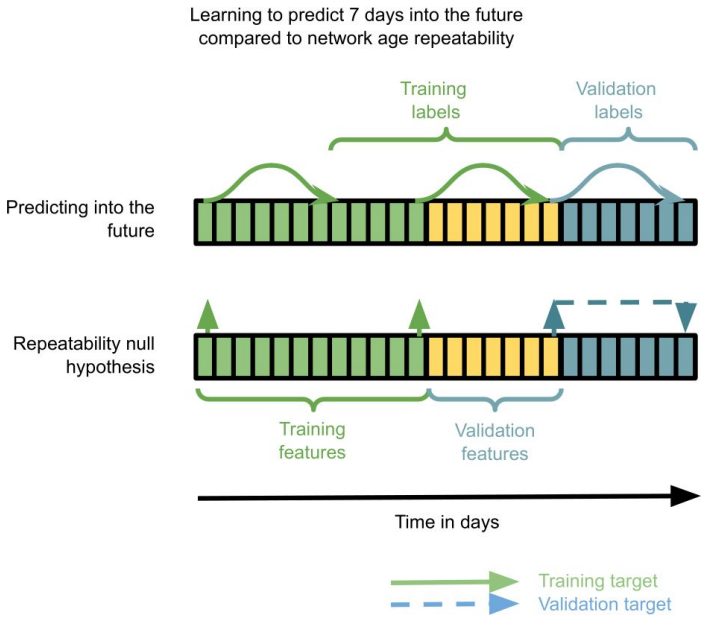

Fig S304C: (Top) Depiction of the process of training a model based on either biological age or network age to predict the future task allocation. We make sure that we do not leak information when evaluating future predictiveness. (Bottom) Data handling for the null hypothesis to check that the predictiveness is not completely explained by the repeatability of network age and task allocation for individual bees.

S305 Probability of dying over the course of a week

We have shown that network age is more predictive of the remaining days to live than biological age (see S303). Here we directly compare biologically young but functionally old bees with biologically old but functionally young bees. To define biologically young and old, we orientate

ourselves at the caste thresholds proposed by Seeley (1982). To define functionally young and old we use the same thresholds for network age as shown in the heatmaps of Fig.3A, because they occupy distinct locations on the comb and therefore hold different functional roles within the colony. However, these exact thresholds are not important as both the effect strength and the significance are high.

We sort the individuals into the two groups based on their network age and biological age. To not count an individual twice, we only consider the last date that a bee was seen in each group. We calculate the probability of dying P as the fraction of bees of each group we do not consider alive after one week (see S002 for more information about the death date of a bee). We calculate a confidence interval by drawing 1000 bootstrap samples for each group.

Biologically young but functionally old bees (biological age < 11 days, network age > 30) have a probability of dying over the course of a week of P=80.6% (bootstrap 95% CI=[74.1%, 87.1%], N=139). Biologically old but functionally young bees (biological age > 20, network age < 15) have a significantly lower probability of dying of P=42.1% (bootstrap 95% CI=[37.4%, 47.2%], N=390; χ^2 test of independence: $p<<0.001$ N=529).