# Modeling Approaches to Understanding Bumblebee Behavior and Population Decline

A Dissertation

Submitted to the Faculty of WORCESTER POLYTECHNIC INSTITUTE

in Partial Fulfillment of the Requirement for the Degree of

Doctor of Philosophy
in
Bioinformatics & Computational Biology
August 2022

**Kevin Nelson Heath**

Approved By:

Dr. Elizabeth Ryder
Professor
Dept. Biology & Biotechnology
Worcester Polytechnic Institute

Dr. Robert Gegear
Assistant Professor
Dept. of Biology
University of Massachusetts Dartmouth

Dr. Dmitry Korkin
Professor
Dept. Computer Science
Worcester Polytechnic Institute

Dr. Sarah D Olson
Professor
Dept. Mathematical Science
Worcester Polytechnic Institute

# Table of Contents

## List of Tables

## List of Figures

## List of Videos

# Abstract

Understanding how anthropogenic disturbances affect plant–pollinator systems at the individual and population level has important implications for the conservation of biodiversity and ecosystem functioning. At the individual level, previous laboratory studies show that anthropogenic disturbances such as pesticides and pathogens, which have been implicated in the rapid global decline of pollinators, can impair behavioral processes needed for pollinators to adaptively exploit floral resources and effectively transfer pollen among plants. However, the potential for sublethal stressor effects on pollinator-plant interactions at the individual level to scale up into changes to the dynamics of plant and pollinator populations at the system level remains unclear. To address this question, we developed an empirically parameterized agent-based model of a bumblebee pollination system called SimBee to test for effects of stressor-induced decreases in the memory capacity and information processing speed of individual foragers on bee abundance, plant diversity, and bee–plant system stability over 20 virtual seasons. Modeling of a simple pollination network of a bumblebee and four co-flowering bee-pollinated plant species indicated that bee decline and plant species extinction events could occur when only 25% of the forager population showed cognitive impairment. Higher percentages of impairment caused 50% bee loss in just five virtual seasons and system-wide extinction events in less than 20 virtual seasons under some conditions. Plant species extinctions occurred regardless of bee population size, indicating that stressor-induced changes to pollinator behavior alone could drive species loss from plant communities. These findings indicate that sublethal stressor effects on pollinator behavioral mechanisms, although seemingly insignificant at the level of individuals, have the cumulative potential in principle to degrade plant–pollinator species interactions at the system level.

Understanding the mechanisms behind bumblebee behavioral response to change is key to improving our ability to model and predict how pollinators respond to rapid human-induced change. Building on the work done with the SimBee model, we implement models of individual memory and decision-making to test behavioral response to simulated scenarios of rapid change. Characterizing the behavioral response to change variability of environments, probability of reward, and frequency of change provide insight into the role of memory and recency effects in adaptive decision-making. In environments with variation, memory provides an adaptive advantage to foraging bumblebees and models of decision-making that utilize memory outperform memory-less strategies. Our tests indicate that recency bias is a possible mechanism that allows bumblebees to adaptively respond to changing and variable environments when new information must be acted upon quickly.

While we establish a foundation for exploring and modeling bumblebee behavior and decisions in plant-pollinator systems at the individual level, improving data collection on the dynamics of plant-pollinator interactions at the population level is critical for conservation efforts. Since long-term controlled experimental studies are difficult to execute, we utilize the citizen science Beecology project to lay foundational work for the automated classification of bumblebee behaviors in videos. Recent advances in deep learning have made rapid and accurate behavior classification of human behaviors possible, but these advances have not been applied to bumblebees. We address this by first creating a dataset of bumblebee action video clips using videos submitted by citizen scientists. The dataset was then used to train and test a two-stream convolutional network (TSN) to test the viability of using deep learning techniques for automated bumblebee behavioral classification. Our work highlights the need for a more robust dataset that can facilitate the use of deep learning architectures.

# Acknowledgements

First, I would like to thank my sources of funding throughout my PhD. This includes funding through NSF Grants 1547291 and 1742446 and the WPI CEDAR fellowship.

I want to thank my advisors Elizabeth Ryder and Robert Gegear for all their help and guidance throughout my Ph.D. They guided me through the difficulties and challenges of interdisciplinary research and taught me how to be an independent researcher. I learned how to approach complex interdisciplinary problems from multiple angles and the true importance of clear communication.

Additionally, I want to thank Kuldeep Gill for the collaborative work on the NSF EARS project. I had a blast learning about the challenges of vehicular communication and the surprising number of parallels between vehicles and bumblebees.

I also want to thank friends and family for being there for me during my PhD. I am grateful to my parents for their continued love and support and making sure I had the education and tools to pursue my dreams. I wouldn't be here without you! I also want to thank Samantha Bryce in particular, as I could not have done this without her unceasing support and patience. Finally, a huge thanks to everyone else who has helped and encouraged me along the way, from my committee members to other graduate students and anyone else I have not yet mentioned.

# Preface

Growing up in Phoenix, Arizona, my parents started a tradition of driving up to Sedona for Thanksgiving. Situated further North and higher in elevation, the air in Sedona would be cool and crisp, a welcome change from the heat of Phoenix. Southwest of the striking Cathedral Rock formation along Oak creek is Crescent Moon Ranch, a property maintained as a historic site by The Forest Service. The water wheel sits unmoving next to a field of dry, brown grass under the clear blue sky, lending a sense of timelessness that leaves one wondering if the scene would look different ten years ago or ten years hence. The creek tells a different story, as leaves change colors and fall into the cold rushing water, carrying away detritus and revealing an unobstructed view of Cathedral Rock. There is no official path over Oak creek there – the last concrete crossing having been washed away long ago – but locals often build stone stepping-stones to replicate the original function of red rock crossing. Some years the stepping-stones stand strong against the current, letting visitors cross the creek without getting a toe wet, while other years the creek swallows the rocks and makes crossing a risky adventure. For me, the yearly visits provided a snapshot of a place that seemed frozen in time. The sky was always clear and blue, the grass always brown and dry, the creek always strong and cold. Yet, I could see the change that was always happening, from the falling leaves to the new stepping-stones every year. No single photo can capture this sense of change, but many photos taken year after year can transform static information into a story of dynamic change.

If I look at a photo of Cathedral Rock from Crescent Moon Ranch, I know far more information than a single photo contains. I know that the trees and grass change over the seasons, I know the creek rises and falls, I know the insects and animals grow and reproduce, and I know the visitors come and go. Studying a location for long enough lets us extract enough details and patterns to imagine how it changes when we are not there to observe it as well as what it may look like in the future. A timelapse – a sequence of photos taken from the same location – can approximate this knowledge in a way that can be shared with others who have never witnessed it themselves. By breaking down a constantly changing and dynamic location into discrete snapshots, we can summarize the details and patterns we witness in a way that can be shared with others. But what if we want to share more than a single location? How can we share a large area, a dynamic experience, or a changing phenomenon?

The details and patterns that describe how something changes, what something looked like in the past, and what we think it will look like in the future can all be recorded and documented in a construct called a model. Models are a way we can share in-depth knowledge of complex events with those who have never experienced them. Like a series of pictures, models are simpler than what they represent but capture the key details and patterns that others need to understand what they are looking at. Despite approximations and idealizations that are pale imitations of the true experience, models let us share important patterns and predictions of the past and future in ways that let other people understand and relate to a phenomenon that they would otherwise not comprehend.

# Introduction

Global bumblebee decline is a phenomenon that needs to be understood and shared. Countless researchers have worked to determine the details and patterns of bumblebee decline, from the way stressors like pesticides and disease affect bumblebees to ways land use can promote or handicap local pollinators. Despite the extensive body of research on bumblebees, existing models have proven to be insufficient to adequately understand and communicate the phenomenon as global bumblebee decline continues[1]. The work described here uses a three-pronged approach to advance efforts to understand, model, and communicate bumblebee decline with a multi-faceted and interdisciplinary approach.

The first prong of this research seeks to understand how anthropogenic disturbances affecting plant-pollinator systems at the individual level can result in population-level decline. In particular, previous laboratory studies have established that pesticides and pathogens that do not out-right kill pollinators can impair behavioral processes needed for them to adaptively exploit floral resources. While these stressors are implicated in global pollinator decline, it is unclear how these sublethal stressor effects on pollinator-plant interactions at the individual level can scale up to alter the dynamics of wild plant and pollinator populations. To address this question, we develop an agent-based model called SimBee to simulate the effects of pesticides and pathogens on bumblebee-plant interactions at the individual level. These stressors are known to impair memory and processing speed in laboratory studies, so we designed and empirically parameterized SimBee to reproduce the effects of those impairments. The model then scales up those individual effects into population-level patterns over relevant temporal scales. We use SimBee to test several scenarios regarding the effects of stressor-induced decreases in memory capacity and processing speed on bee and plant abundance. These findings demonstrate that sublethal stressor effects on individual pollinators and plants can have a cumulative effect on plant-pollinator interactions at the population level. Agent-based models are best-suited for capturing these individual-level details in ways that can explain population-level patterns.

The second prong builds from the SimBee model developed previously to investigate and model how bumblebees respond to change. Bees can encounter several different types of change in their environment, and better characterizing how they use memory and decision-making skills to respond to these changes is important for understanding how bees can adapt to human-induced change and furthering our ability to model bees at the individual level. Specifically, we explore how bumblebees respond and adapt to a change in the variation of resources, change in the probability of reward, and changes in the frequency of change in their environment. While the underlying role of memory in response to change – as well as variation – are not well understood in bumblebees, the topic has been explored in human psychology. We design three scenarios based on experiments from human psychology for investigating bumblebee response to change in three different ways and simulate results using the SimBee model. The results demonstrate that memory and complex decision-making strategies can provide adaptive benefits in some scenarios, and the tested models provide insight into future directions for studying the bumblebee response to change.

Finally, the third prong explores deep learning modeling techniques to address current challenges with collecting behavioral information for conservation efforts. Deep learning models for human action classification are continually advancing and demonstrate a high degree of accuracy when tested on a diverse range of human behaviors. Capturing bumblebee behaviors from videos using deep learning models would greatly benefit conservation efforts since the analysis of publicly submitted videos is a serious bottleneck in executing large-scale studies of pollinator decline. While honeybees (*Apis mellifera*) are managed for agricultural purposes, bumblebees (genus *Bombus*) are wild native pollinators in the United States and difficult to document on a large scale. However, deep learning models for action classification are rarely applied to non-human organisms and there is currently no open-source dataset or behavior classification system for bumblebees. The Beecology Project[2] is a citizen-science driven project that aims to collect ecological data on native pollinators, and user-submitted videos enable the creation of a dataset for training deep learning models. In this work, we use bumblebee video data to train and test an established two-stream convolutional network model for bumblebee behavior classification. The results show that advances in human behavior classification can be applied to bumblebee behaviors and provide insight for future work on this topic.

## Bumblebees

Bumblebees are members of the genus *Bombus*, which includes approximately 260 wild species spread across most of the world, 47 of which are found in North America[1]. Most bumblebees are eusocial insects that form colonies with a single queen, but some species are solitary[3]. They are round and covered in soft hair that gives them a 'fuzzy' appearance[4]. Bumblebees act as natural pollinators while they gather nectar and pollen from flowering plants, and their use of buzz-pollination - a type of vibration that helps shake pollen from flowers – is key to the efficient pollination of plants such as berries and tomatoes[5]. While bumblebees are commonly confused with the western honeybee (*Apis mellifera*) by the general public, honeybees and bumblebees are different in numerous ways and research on one cannot necessarily be generalized to the other[4].

When studying bumblebee populations, it is important to note that bumblebees are annual social insects and have a distinctly different lifecycle than honeybees (see Figure 1 for a summary of the bumblebee lifecycle)[6]. Queens emerge from hibernation in the spring and found individual colonies. Once the first workers hatch, they help establish the colony and forage for nectar and pollen to support population growth. Towards the end of the colony cycle in the late summer, the colony transitions from producing workers to producing reproductives (males and gynes). After mating, the young queens go into hibernation while the founding queen, workers, and males all die. Queens that survive hibernation give rise the next generation in the following spring[7].

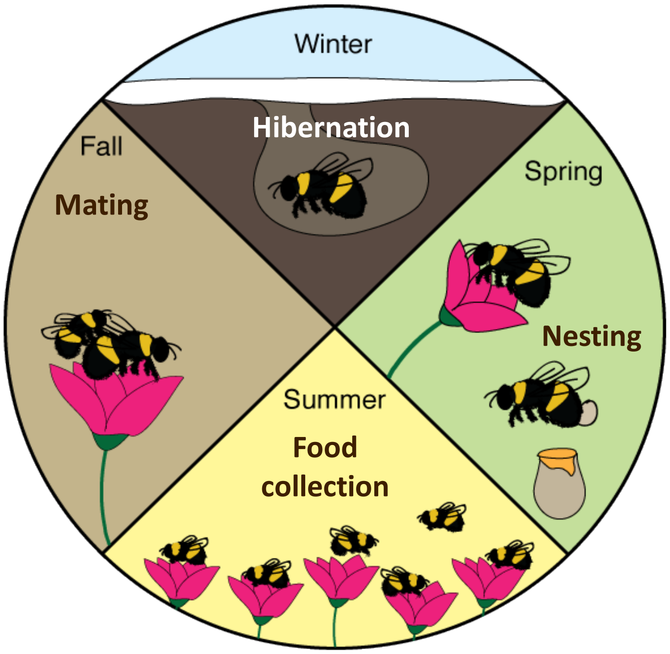

*Figure 1: Summary of the bumblebee lifecycle. Credit to Rachel Blakely for the image.*

Due to the annual lifecycle, bumblebees do not remain in the same location from year to year nor do they maintain large stores of honey like honeybee colonies[8]. While some of these differences make bumblebees a poor choice for modern agricultural practices, many species of bumblebees are incredibly important pollinators for key crops[5,9,10]. In fact, certain plants such as tomatoes and cranberries rely almost exclusively on bumblebees for pollination and cannot be efficiently pollinated by honeybees[9]. Bumblebee visitation also enhances the yield of many other crops[11].

Additionally, most wild plants are pollinated by different species of bumblebees[5]. Due to the interconnected nature of pollinator-plant networks, eliminating bumblebee species from pollination networks likely leads to serious decline with respect to plant diversity[12,13]. Thus, the conservation of bumblebee species is important for biodiversity and economic reasons.

## Bumblebee Decline

Over the past decade, wild pollinators have declined in abundance, species richness, and geographic distribution at an alarming rate worldwide[14,15]. A great deal of progress has been made in evaluating the extinction risk of bumblebee species, but there are still many unknowns regarding the distribution and decline of different species in various regions (See Figure 2). Although the cause of these declines is unknown, human-introduced stressors such as pesticides, disease, habitat loss, and climate change have all been identified as potential contributing factors[11]. Given that many wildlife species depend on animal-pollinated plants for food, shelter, and nesting habitat, pollinator loss has the potential to significantly degrade the function and diversity of terrestrial ecosystems. Indeed, abundance and diversity of flowering plants have declined in parallel with their pollinators in many locations[16], suggesting that these cascading negative effects might be well underway. A critical step in the conservation and restoration of wild pollination systems is understanding how anthropogenic stressors affect the dynamics of plant-pollinator interactions over different levels of biological organization and over ecologically relevant temporal scales. However, such data are often difficult to obtain due to the significant logistical challenges associated with conducting long-term controlled experimental studies of natural systems[17,18].

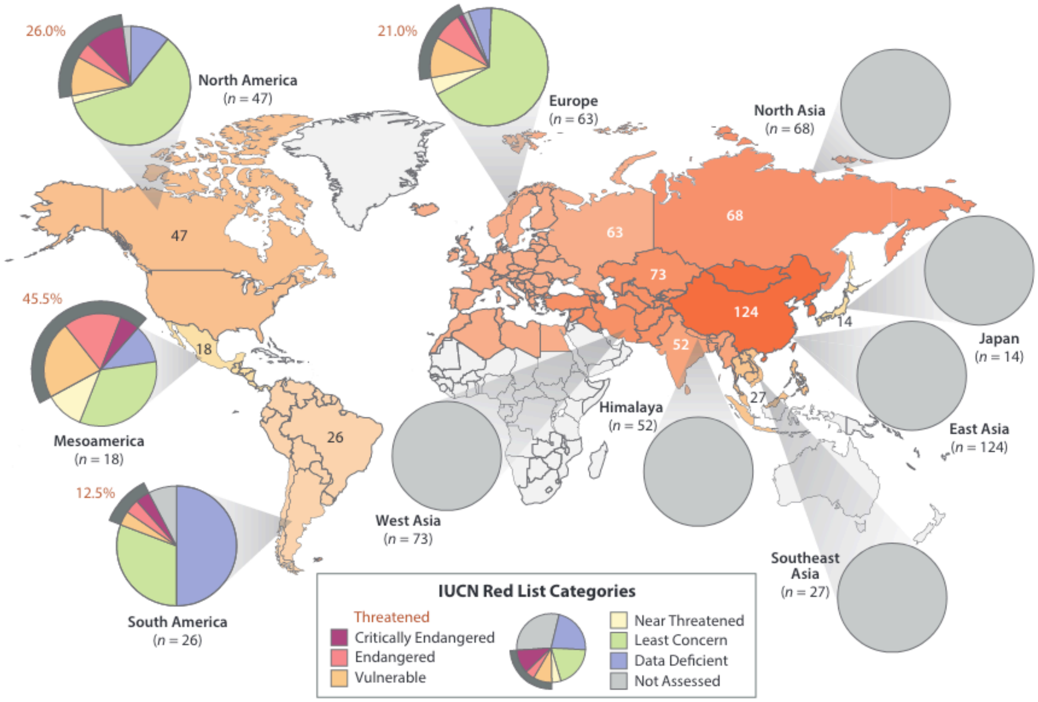

*Figure 2: Global map of International Union for Conservation of Nature (IUCN) Bumblebee Specialist Group regions (color-coded in shades of red), each displaying a pie chart indicating proportions of the different Red List threat categories assessed for the bumblebee species of a given region. Percentages refer to the fraction of IUCN-assessed species designated as Threatened. At time of publication, 154 species total have been described for Europe, North America, Mesoamerica, and South America, and 150 species have been IUCN assessed, of which 36 (24%) are currently listed as threatened. Note that regional species totals are not mutually exclusive since some species occupy multiple regions. Taken from Global Trends in Bumble Bee Health*[1].

## Sublethal Stressor Effects of Pesticides & Pathogens

Neonicotinoid-based pesticides were originally assumed to be toxic to pest insects but relatively harmless towards pollinating insects such as honeybees and bumblebees[19–22]. This was supported by studies showing that field-realistic exposure of bees to neonicotinoids in nectar and pollen of seed-treated crops is unlikely to cause lethal effects[23]. While neonicotinoids are unlikely to cause direct mortality in the field, numerous studies in the past decade have demonstrated that neonicotinoid use has indirect (sublethal) effects on bees[24–27].

A sublethal effect is generally defined as a stressor-induced change to an individual's physiology or behavior that does not cause direct mortality[25,28]. Neonicotinoids work by binding to nicotinic acetylcholine receptors in the synaptic membrane of insect neurons, resulting in overactivation and the eventual death of the neuron[1]. There is a substantial amount of empirical evidence that chronic exposure to neonicotinoids can decrease foraging success in bee pollinators through impairments to underlying cognitive mechanisms[27,29–33,33–37]. For example, Stanley et al. (2015)[38] showed that 50% of bumblebees (*B. terrestris*) could not remember the reward properties of available flowers 24 hours after chronic exposure to a field-realistic dose of a neonicotinoid pesticide. However, pesticides are not the only source of sublethal stressors for bumblebees. The prevalence and diversity of bumblebee pathogens has greatly increased over the past decade, and the increased expose to pathogenic organisms is thought to be a contributing factor to bumblebee decline[39,40]. In another bumblebee species (*B. impatiens*), Mobley and Gegear (2018)[39] found that individuals with an activated immune system took 40% longer to make foraging decisions than controls, significantly reducing the rate of reward delivery to the colony. Numerous other studies have demonstrated that sublethal effects can impair the ability of bees to forage, learn and remember flower locations, and negatively impact the growth of colonies[23,41–50]. Despite the growing body of evidence that the sublethal effects of pesticides and pathogens have a serious long-term impact on the cognitive abilities and reproduction of bumblebees, the indirect nature of sublethal effects makes it unclear if such changes to cognition-based behavior at the individual level are sufficient to drive the observed global population decline. While a recent large-scale study in avian species concluded that the increase in neonicotinoid use led to statistically significant reductions in bird biodiversity[51], there is no similar large-scale study of bumblebee biodiversity. A critical step in the conservation and restoration of wild pollination systems is understanding how anthropogenic stressors affect the dynamics of plant-pollinator interactions over different levels of biological organization and over ecologically relevant temporal scales. However, such data are often difficult to obtain due to the significant logistical challenges associated with conducting long-term controlled experimental studies of natural systems[17,18].

## Pollinator-Plant Dynamics

The long history of coevolution between flowering plants and pollinators such as bumblebees has resulted in a complex web of interactions that make both partners dependent on one another[52–55]. This partnership between plants and pollinators is based on mutualistic interactions driven by selfish interest[16,54–56]. Plants provide nectar and pollen to reward visiting pollinators, who in return sustain plant reproduction by vectoring pollen to conspecific flowers[52,55]. From the perspective of a pollinator, a field full of flowers is a rich marketplace where numerous floral species offer nectar and pollen rewards. However, foraging for these rewards is not a trivial task.

In the context of bumblebees and plant-pollinator interactions, there are several challenges that foraging pollinators face while searching for resources. First, foragers may encounter spatial variation in the location of resources in the local environment;[57–61] foragers may find no resources in one area or multiple resources in another. Second, foragers also experience temporal variation in resources. For example, plants go in and out of bloom at different times during the year and this also changes what resources are available to foragers[53,53,62,63]. These changes fall under phenology, the periodically recurring patterns of growth and development of plants and animal behavior during the year[64]. Finally, bumblebees must navigate these challenges to determine which of the many available resources is most rewarding. Limits of time and energy require foragers to make economic decisions to ensure a net gain of resources[52]. In bumblebees, resource collection is strongly correlated with colony fitness[65,66]. Other factors include the ease of handling, the presence of distractions, and competition from other pollinators. Foragers must consider all these variables to successfully utilize the floral marketplace.

From the plant's perspective, pollinators are greedy customers that unwittingly become biotic gene dispersers. Pollinators that stick to the same floral species assist plant reproduction while pollinators that switch species can impair the reproductive potential of the plants they visit[54,59,67,68]. In order to attract and retain these customers, flowering plants have evolved a number of traits to attract or deter specific pollinators[52,55,61,69]. While visual displays, olfactory cues, and physical shape certainly play an important role, the most important strategy plants employ is providing the best nectar and pollen rewards in the market[61,70–73]. Consequently, visiting foragers that were initially attracted by visual or olfactory stimulus learn to associate these cues with the nectar and pollen rewards provided.

## Memory in Bumblebees

While the complex behaviors of groups of eusocial insects have long been recognized, the behavioral complexity of the individuals has often been ignored or dismissed. The study of cognition in insects broadly seeks to understand the mechanisms by which insects acquire, process, store, and act on external stimuli[74–76]. Tests for associative learning, sensory integration, and choice behavior are the most common methods for determining the presence of cognition, as higher level capabilities such as concept learning, cognitive maps, and metacognition can be difficult to demonstrate in insects[74,75,77,78].

Eusocial bees of the *Apidae* family are excellent systems for studying how learning and memory occurs and its ecological consequences for a multitude of reasons. Bees forage in complex and dynamic environments where they regularly visit plants that vary in their floral traits and reward composition. Individual foragers must assess observed and experienced differences, as well as rapidly form associations between floral features and rewards[54,71,79–84]. Additionally, both individual foragers and colonies vary in their stimuli (odor and color) preferences and cognitive abilities[66,82,85–87]. These variations in preferences and learning have implications for colony-level fitness as well as the fitness of mutualistic partners. Rapid learning of floral cues is strongly correlated with improved foraging efficiency[66]. In a study on colony learning speed, the slowest learning colonies brought in 40% less nectar than the fastest learning colonies[66]. Suboptimal learning may lead to pollen transfer between two different plant species that incurs a reproductive cost to the receiving plant (heterospecific pollen transfer)[66,66,82,88]. Unfortunately, learning and memory can be challenging to quantify in field populations since restraining and testing individual bumblebees is nearly impossible in the wild. Experiments must be carefully designed to properly measure interspecific or intraspecific variation in cognitive performance[89–91]. Due to numerous confounding variables, the vast majority of experiments on cognition occur in lab settings on three commercially available social species: the honeybee (*Apis mellifera*) and the bumblebees (*Bombus impatiens* and *Bombus terrestris*)[92]. Unfortunately, most of these studies focus on *A. mellifera* and there are large gaps in knowledge when it comes to cognition in bumblebee species. Modeling bumblebee memory and decision-making is paramount for addressing this knowledge gap in understanding how and why bumblebees respond to changing stimuli.

The ability of organisms to recall past information/associations and apply it to decision-making is generally attributed to some form of working memory. Working memory refers to the temporary storage of information in connection with the performance of other cognitive tasks[93]. The size and duration of working memory is considered a crucial component of intelligence in mammals and insects[75,94–96]. There exist two primary paradigms for testing working memory in bees: 1) the conditioning of the proboscis extension response (PER) in restrained bees and (2) the training of the approach flight toward a visual target in free-flying bees. In PER, harnessed bees are trained to associate a floral trait with a sucrose reward – the extension of the proboscis indicates that the bee has 'learned' the association[97–99]. In flight tests, free-flying bees are pretrained to fly toward the training/test destination by associating the presence of a visual clue with a sucrose reward – both decision-making and memory can be tested by varying the location of correct and incorrect visual clues[95,99]. The results from experimental studies using these paradigms are important for understanding how bees learn and maintain information over short periods of time. However, there is a gap between experimental studies of and models of bumblebee working memory and decision-making.

## Agent-Based Modeling

Agent-based modeling (ABM) is a powerful and widely used tool for simulating complex systems of autonomous agents in a variety of disciplines, such as ecology[100–102], social sciences[103], and economics[104]. ABMs are widely used in ecology as a complementary approach to conducting studies of complex natural biological systems such as those involving interactions between predators and prey[105] or pollinators and plants[106–108]. Unlike traditional differential equation-based models, ABMs are designed to investigate how the impact of individual variation in particular characteristics can influence processes at greater levels of organization. The model functions as a system of behavioral rules that model the biology of individuals utilizing information from lab and field studies. The results of interactions between individuals with unique histories aggregate at the population level to produce novel system patterns. While the concept of ABMs is not new, the recent surge in popularity over the past decade has produced a wide variety of models and approaches.

Eusocial bumblebees – species with an advanced level of social organization where a single female or caste produces offspring and non-reproductive individuals care for the young - are especially suitable for ABMs due to their independent foraging behavior[17,109–111]. While honeybee scouts and foragers communicate with each other about sites where nectar and pollen are available using the well-known dance language, bumblebees do not[92]. Although bumblebees can obtain information that floral resources are available from the pheromone signals of foragers returning to the nest, they ultimately make independent decisions in the field based on their individual experience of the floral environment[8,112–115]. This independent decision-making makes individual bumblebee foragers ideal candidates for ABM.

## Models in Bumblebee Research

There are several examples of ABMs in bumblebee research that use individual parameters and behaviors to investigate population level patterns. One extremely relevant example is BEE-STEWARD's[116] Bumble-BEEHAVE[100] model component, an agent-based systems model for simulating bumblebee populations in agricultural settings. While it incorporates numerous aspects of pollinator-plant interactions, the general-purpose nature of the model makes it difficult to ask in-depth questions about the sub-lethal effects of pesticides and pathogens on working memory. The Bumble-BEEHAVE model simulates the effect of colony-level pesticide exposure with reproduction depression, killing 26% of queens that emerge from hibernation each year[100]. Their approach uses cohorts of foragers with a single memory and decision-making process for the entire cohort. It also does not address the sub-lethal effects of pesticides on memory, nor is it possible to investigate the effects of partial colony exposure to pesticides. Another agent-based model - EcoSimInGrid - implements adaptive memory for pollinators but does not include cognitive impairment from stressors[15,27,31,33,33,117]. It is also primarily focused on plant population dynamics and does not include the effect of pollinator-plant interactions on pollinator population dynamics[109]. BeeNestABM is an agent-based model of bumblebee that models the spread of pesticides within a nest[118]. This model uses bumblebee activity and movement data to simulate how the behaviors of bees exposed to pesticides affect overall colony fitness. Beyond the agent-based models described here, there also other equation-based models of note.

There are several existing models that take a traditional mathematical approach to bumblebee population modeling. The earliest model covered here, Bryden et al. (2013), used a system of differential equations to model the sublethal effects of pesticides on bumblebee colony size[119]. Their model incorporated sublethal effects by increasing the mortality rate of impaired bees. Cresswell (2017)[23] uses a Markov matrix model to simulate the effects of pesticides and predation on bumblebee colonies. These simulations were compared to existing studies to determine if a reduction in workers due to pesticides could explain observed population decline.

Next, the Banks et al. (2017) model established a much more detailed delay differential equation (DDE) population model of bumblebee colonies[120]. The model incorporates both resource abundance (nectar and pollen) as well as the changing populations of queens, workers, males and gynes in the colony. Although this DDE model did not feature the effects of pesticides, it served as the foundation for their next model (Banks et al. 2020[121]) specifically designed to model the sublethal effects of pesticides on populations. This more recent model adds larval development as well as a combination of lethal pesticide effects on workers and sublethal effects via reduced foraging abilities and brood sizes. Although the DDE model clearly demonstrates that mortality and reduced resources due to pesticides negatively affect colony size, it lacks mechanistic explanations for how sublethal effects result in reduced resource intake. Overall, an individual-level approach is needed to explore why sublethal pesticide impairments may negatively impact bumblebee populations over longer temporal scales.

Other bumblebee-focused models simulate the effects of landscape and floral resources on bumblebee populations. Both Olsson et al. (2015) and Haussler et al. (2017) aim to predict the effects of landscape use on bumblebee flower visitation rates. These models are especially useful for farmers and agricultural studies to determine land-use practices that benefit wild bumblebees. Crone and Williams (2016)[111] aims to predict the impact of floral resources on colony growth and makes the case that increasing the availability of floral resources does increase queen production in real landscapes. These models are important for conservation efforts but do not address the effects of pesticides on bumblebee decline. See Table 1 for a detailed comparison of the models mentioned in this section.

## Details of Bumblebee models

*Table 1: Comparison of the aims, processes, and output captured by different bumblebee models. '+' indicates something is explicitly included in the model, (+) means only implicitly included or the authors state that this could be simulated.*

| ***Comparator*** | **Biernaskie et al. (2009)[122]** | **Qu et al. (2013)[109]** | **Bryden et al. (2013)[119]** | **Olsson et al. (2015)[110]** | **Crone & Williams (2016)[111]** | **Cresswell (2017)[23]** | **Banks et al. (2017)[120]** | **Haussler et al. (2017)[123]** | **Becher et al. (2018)[100]** | **Versypt et al. (2018)[118]** | **Banks et al. 2020** |
|---|---|---|---|---|---|---|---|---|---|---|---|
| *Model Name* | Bumblebees Learn to Forage Like Bayesians | EcoSimInGrid | SLS Model | Central Place Foraging (CPF) model | Colony Growth model | Demographic Matrix model | DDE Population Dynamics model | Pollinator Land-Management model | Bumble-BEEHAVE | BeeNestABM | DDE Toxicology Population model |
| *Type of model* | Optimal Foraging model | ABM, Reinforcement Learning | Differential Equations | Distance decay | Statistical, Differential equations | Matrix based demographic models | Delay Differential Equations | Process-based | ABM, Monte Carlo | ABM | Delay Differential Equations |
| *Model aims to predict:* | When bumblebees leave a patch of resources | Effects of shared pollination services on plant communities | Impact of sublethal stress | Effect of landscape on flower visit rate and bee fitness | Impact of floral resources on colony growth and queen production | Colony demography and impact of pesticide and predation | Impact of many stressors on multiple colony growth | Effect of landscape on flower visit rate | Impact of many stressors on individual, colony, population, and community - with mapping | The localized responses of bumblebees to sublethal exposures of neonicotinoids in the hive. | Lethal and sublethal effects of pesticides affect population outcome |
| *Main outputs* | Giving-up times on exhausted patches | Flower visit rates, flower success and diversity, pollinator behavior | Colony size, survival | Nest fitness, flower visit rates | Colony size and mass, queen productions, survival | Colony size, reproduction, survival | Colony size and composition, queen and male production, stores, survival | Number of colonies, survival, flower visit rates | Behavior, number of colonies, size and composition, stores, queens, and males, production, survival, flower visits | Spatial location in hive over time | Cumulative reproductive output (males and gynes) |
| *Scale (grid size / map size)* | Infinite series of patches | 500 * 500 cells | 25 m / 3km | | | | | 25 m / 3 km | 25 m / 5 km | 25 cm x 20 cm | |
| *Time* | Seconds | 1 year time steps | Continuous | | Discrete, 15 weeks | 1 day steps, 40 days | Continuous, 120 days | | | 2 Hz time step, 5-60 minutes total | 1 day steps, 100 days |
| *Environment* | Patches of flowers | Agricultural field | | | | | | | | Bumblebee Nest | |
| | | | | | | | | | | | |
| *Organizational level:* | | | | | | | | | | | |
| *Individual Behavior* | + | + | (+) | | | | | | + (queens only) | + | |
| *Energy/nectar consumption* | + | | | + | | | + | | + | | + |
| *Flower pollination* | | + | | | | | | | | | |
| *Colony level* | | | (w) | | Colony founding queens, offspring queens, males; mass | Workers, males, offspring queens | Colony founding queens, larvae, workers, males, offspring queens; nectar and pollen stores | Colony founding queens, offspring queens, workers | Colony founding queens, eggs, larvae, pupae, workers, males, offspring queens; nectar, pollen | | First brood of workers, larvae, males and gynes. |

| | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| *Flower lifecycle* | | + | | | | | | + | + | | |
| *Bee lifecycle* | | | | | | | + | | + | | + |
| *Foraging behavior* | + | + | | | | | | | + | | + |
| | | | | | | | | | | | |
| *Stressors:* | | | | | | | | | | | |
| *Foraging availability* | | + | | + | + | | + | + | + | | |
| *Nest site availability* | | | | + | | | | + | + | | |
| *Pathogens/Parasites* | | | (+) | | | | | | (+) | | |
| *Predation* | | | | (+) | | + | | | + | | |
| *Pesticides (lethal)* | | | + | | | + | | | + | | + |
| *Pesticides (sub-lethal)* | | | | | | | | | | + | + |
| *Weather/Climate* | | | | | | | | | (+) | | |
| *Competition* | | | | | | | | | + | | |
| | | | | | | | | | | | |
| *Testing* | | | | | | | | | | | |
| *Sensitivity analysis* | | | | | + | | + | | + | | + |
| *Real-world comparison/verification* | + | + | + | | + | + | (+) | (+) | + | | |
| *Memory Model* | Bees remember past rewards | Reinforcement Learning | No memory | Bees don't make decisions | Bees don't make decisions | Bees don't make decisions | | | Bees remember every patch they visit. Foragers choose based on time to collect a pollen load or max energetic efficiency for nectar. | Bee movement is determined by bee density. | No individual level behavior |

## Using Citizen Science to Improve Conservation Efforts

Bumblebee population decline is just one of many current environmental challenges of enormous scale and complexity for the field of conservation biology. Understanding how anthropogenic stressors affect the dynamics of plant-pollinator interactions over different levels of biological organization and over ecologically relevant temporal scales is critical for the conservation and restoration of wild pollination systems[13,5]. While long-term controlled experimental studies are difficult to execute, citizen science provides an alternative approach to gathering such data and disseminating scientific knowledge to the public.

Although citizen science is often used in many different contexts and meanings, in research terminology it represents the practice of engaging the public in a scientific project to produce reliable data and information[124]. Citizen science projects typically involve a large number of people involved in the data collection driven by scientific questions and hypotheses developed by professional scientists Unlike experimental studies, citizen science data is often easier to collect than to validate[125,126]. Despite this, recent studies have demonstrated that citizen science can improve conservation efforts by building scientific knowledge, informing policy formulation, and inspiring public action[124]. Indeed, the nature of citizen science makes public engagement a key component of such projects in ways that traditional methods of scientific dissemination have difficulty replicating. Due to the value of citizen-gathered data and engagement, citizen science has become a key component in understanding and combating ecological change.

An example of an ongoing citizen science project is the Beecology Project[2]. This project seeks to gather critical ecological information for bumblebee conservation by recruiting citizen scientists to gather and submit data on native pollinator species. The Beecology Project develops and runs an online tool for submitting pollinator information as well as several online visualization tools for educational and conservation-related purposes. Importantly, the online tools allow the general public to submit image and video records of local pollinators to a public database that provides valuable information to researchers and conservationists.

In order to improve the utilization of citizen science data, the field of conservation ecology is in need of better methods for analyzing and validating data submitted by the general public[125–127]. In particular, sustaining the collection of high-quality data can be difficult when the expertise required to properly identify and log information is uncommon in the general population[128]. For example, properly identifying on of the 21 species of bumblebees present on the East coast of the United States generally takes extensive practice, even for professional scientists. While identification guides can help decrease false identifications, errors in data collection persist and must be corrected manually by experts. More complex data - such as observed behaviors in bumblebees – are extremely valuable in conservation research but present equally serious problems with data collection and validation. The relative lack of experts compared to the population of citizens capable of submitting information presents a significant bottleneck in the validation and utilization of publicly gathered data.

# Chapter 2: SimBee ODD

The work presented in this chapter is published in Volume 35, Issue 5 of Conservation Biology as supplementary information and is reproduced below.

Gegear, R. J[*]., Heath, K. N[*]. & Ryder, E. F. Modeling scale up of anthropogenic impacts from individual pollinator behavior to pollination systems. *Conservation Biology* **35**, 1519–1529 (2021).

The model description follows the ODD (Overview, Design concepts, Details) protocol, a standard format for describing individual-based models[129]. SimBee was implemented in NetLogo[130], version 6.0.4. The ODD was written by Kevin Heath and edited by Robert Gegear and Elizabeth Ryder.

## 1. Purpose & Patterns

The purpose of the SimBee model is to explore how sublethal effects of stressors on bumblebees at the individual level 'scale up' to impact the structure and diversity of the bumblebee pollination system. The model uses an individual agent-based approach to simulate wild bumblebees foraging for nectar and pollinating plants in a spatially explicit landscape. Starting with initial populations of bumblebees and plants, we model the behavior of individual foragers and the effect they have on plant pollination over multiple seasons. Population dynamics for both bumblebees and plants result from the amount of nectar reward collected over a season and the

---

[*] These authors contributed equally to this work

number of seeds produced for each species of plant. The model interface is shown below (Figure 3).

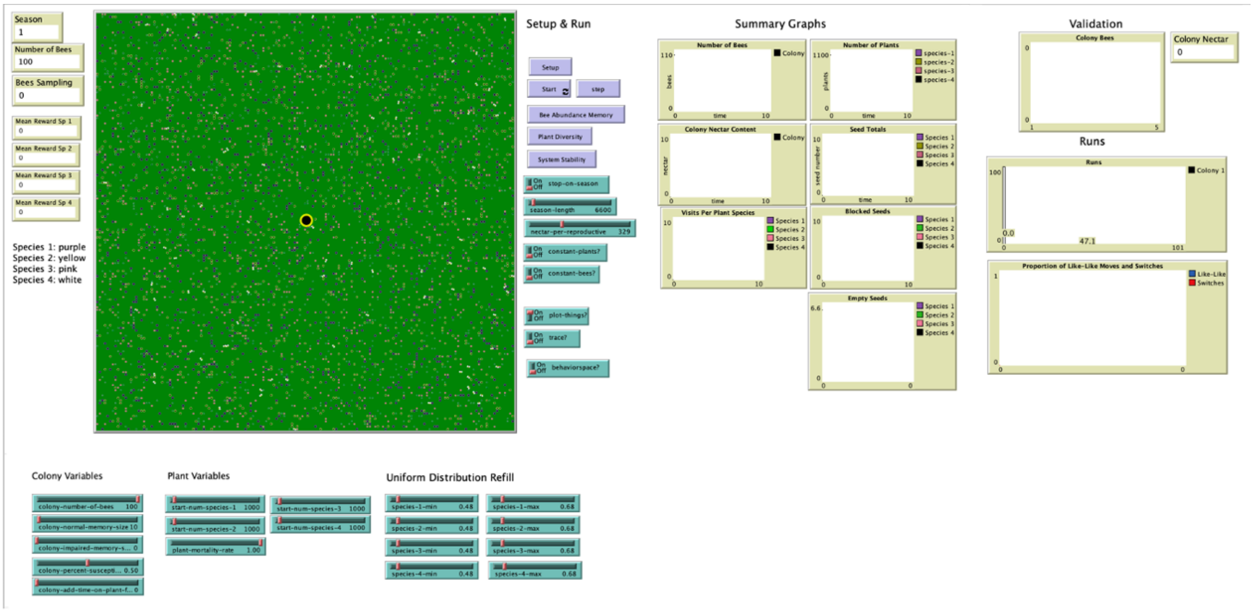

Season 1
Number of Bees 100
Bees Sampling 0
Mean Reward Sp 1 0
Mean Reward Sp 2 0
Mean Reward Sp 3 0
Mean Reward Sp 4 0

Species 1: purple
Species 2: yellow
Species 3: pink
Species 4: white

season-length 6600
nectar-per-reproductive 329

Colony Nectar 0

Colony Variables
colony-number-of-bees 100
colony-normal-memory-size 10
colony-impaired-memory-s... 0
colony-percent-suscepti... 0.50
colony-add-time-on-plant-f... 0

Plant Variables
start-num-species-1 1000
start-num-species-2 1000
start-num-species-3 1000
start-num-species-4 1000
plant-mortality-rate 1.00

Uniform Distribution Refill
species-1-min 0.48
species-1-max 0.68
species-2-min 0.48
species-2-max 0.68
species-3-min 0.48
species-3-max 0.68
species-4-min 0.48
species-4-max 0.68

*Figure 3: Screenshot of the model interface. On the right are plots that display information about agents in the simulation. On the top left are trackers for the most common elements of the model. The green area is a spatial representation of the environment. Controls for setting up and running the simulation are located to the right of the environment. Below the environment are sliders that control variables in the model that should be specified by the user.*

To show that our model is realistic enough for its purpose, we evaluated its ability to reproduce three key patterns found in natural systems (see Chapter 3: Modeling scale up of anthropogenic impacts from individual pollinator behavior to pollination systems Results, and Appendix S1: Validation of Model Assumptions):

Pattern 1: Working memory should enable bees to adaptively exploit flowers under highly variable reward conditions.

Pattern 2: Plant reproductive success should depend on the type of pollen transferred among flowers (determined by flower choice behavior of bees) and the amount of pollen transferred (determined by the number of bees in the system).

Pattern 3: Bees should exhibit 'floral constancy', which is the propensity to make 'like-like moves', that is, visit flowers of the same species of plant in sequence[67].

## 2. Entities, State Variables, and Scales

The SimBee model has three agents: the *bee*, the *plant*, and the *colony*. *Bee* agents actively explore the environment and visit *plants* to collect nectar and pollen. *Plant* agents are stationary and keep track of seeds that are produced after pollination by visiting *bees*. The *colony* agent holds the nectar that *bee* agents collect and produces new *bees* based on that amount each season.

### A. Bumblebees

We implemented bumblebees using the NetLogo *breed [bees bee]*. Each *bee* represents a single adult forager. *Bees* primarily vary in memory (how many past visits they can recall) as well as several other individual traits. Additionally, *bees* belong to a species defined by the *colony* they belong to. There is only one colony in this version of the model, but the framework for two colonies is included. The relevant variables for the *bee* agent can be found in Table 2.

### B. Plants

*Plants* are implemented in NetLogo as *breed [plants plant]*. There can be no more than a single *plant* in each *patch*. Every *plant* belongs to one of four species that differ in color. *Plant* species can vary in reward. The SimBee model uses uniform random distributions to determine the reward a *plant* contains when it is visited. See Table 3 for specific variables and values. In addition, every *plant* starts the season with a set number of potential seeds that can only be fertilized by pollen (detailed later) from the same species. Pollen from a different species blocks the fertilization of a potential seed.

## C. Colony

*Colonies* are implemented as NetLogo *breed [colonies colony]*. Each *colony* can represent either a single colony of *bees* or a population of *bees* of the same species. The *colony* records the total nectar collected by *bees* over the course of a season. Since the associated paper only tests experiments with a single colony, the multiple colony feature is disabled for this version of the model. See Table 4 for *colony*-specific variables.

## D. Environment

### *I. Grid Cells*

The NetLogo world (the visual representation of the model landscape) is made of a grid of 161 x 161 cells (referred to as *patches* in NetLogo). The real dimensions of a grid cell vary with the landscape simulated and there is not necessarily a direct translation between cells and meters. See Table 5 for environment and other general variables.

### *II. Time*

Each time unit (tick) is on the order of a second, approximately the amount of time it takes a *bee* to make a decision. This small-scale time interval allows us to capture the benefits and costs of individual decisions. Since our model is focused on individual-level decisions of pollinators, we only model time periods when the pollinators could be foraging. Consequently, we do not implement a day/night cycle or weather that could prevent foraging. Due to the small scale, it is not feasible to run the model for an entire season, so we use 6,600 ticks as a virtual season. The impairments we model at the individual level, decreased memory capacity and processing speed, persist throughout a season, so it is expected that our simulated virtual season captures the impact of the behavior over a full season.

*Table 2: Bee Variables*

| **Bee Variables** | **Description** | **Range** | **Default Value** |
|---|---|---|---|
| bees-start-with-sample? | Should bees immediately sample after leaving the colony at the beginning of the season? | TRUE / FALSE | TRUE |
| constant-bees? | Population of each colony remains constant between seasons | TRUE / FALSE | FALSE |
| delay-max | Maximum number of ticks a bee will wait before leaving its colony at the beginning of the season | 0-200 | 100 |
| nectar-per-reproductive | Amount of nectar required to produce a new bee for the next season | 0-1000 | 181.5 |
| bee-nectar-max | Maximum units of reward a bee can hold before needing to return to the colony | 0-200 | 80 |
| colony-1-add-time-on-plant-for-impaired[1] | Additional number of ticks a bee from colony 1 impaired by pesticides must spend on plants | 0-100 | 0 |
| colony-1-handling-time-species-1 | Number of ticks it takes bees from colony 1 to gather reward from a plant of species 1 | 0-100 | 5 |
| colony-1-handling-time-species-2 | Number of ticks it takes bees from colony 1 to gather reward from a plant of species 2 | 0-100 | 5 |
| colony-1-handling-time-species-3 | Number of ticks it takes bees from colony 1 to gather reward from a plant of species 3 | 0-100 | 5 |
| colony-1-handling-time-species-4 | Number of ticks it takes bees from colony 1 to gather reward from a plant of species 4 | 0-100 | 5 |
| colony-1-impaired-memory-size | Number of past visits a bee from colony 1 impaired by pesticides can remember | 0-20 | 1 |
| colony-1-normal-memory-size | Number of past visits a bee from colony 1 can remember | 0-20 | 10 |
| colony-1-number-of-bees | Initial number of bees in colony 1 | 0-100 | 100 |

---

[1] As mentioned above, only one colony is implemented in this version of the model. Colony-2 is intended for future use.

| | | | |
|---|---|---|---|
| colony-1-percent-susceptible | Proportion of bees in colony 1 susceptible to pesticides | 0-1 | 0 |
| colony-1-pref-increase | Percent increase in preference for bees in colony 1 for a species of plant | 0-1 | 0 |
| colony-1-preference | Species of plant that bees from colony 1 show preference for | 1-4 | None |
| colony-1-close-enough | On a scale from 0 to 1, what percent difference in reward can bees from colony 1 distinguish? | 0-1 | 0.1 |
| colony-1-all-knowing | The bees from colony 1 calculate the mean reward for each species of plant from the past *sliding-window-size* number of visits to each species. They use this value instead of the values in memory to make decisions. | TRUE / FALSE | FALSE |

*Table 3: Plant Variables*

| **Variable Name** | **Description** | **Range** | **Default Value** |
|---|---|---|---|
| constant-plants? | Population of each species of plant remains constant between seasons | TRUE / FALSE | FALSE |
| use-normal-distribution? | If *instant-refill?* is TRUE, should plant rewards refill based on a standard normal distribution (TRUE) or uniform random distribution (FALSE)? | TRUE / FALSE | FALSE |
| species-1-max | Maximum reward value for species 1 when drawn from a uniform distribution | 0-5 | 0.5 |
| species-2-max | Maximum reward value for species 2 when drawn from a uniform distribution | 0-5 | 0.48 |
| species-3-max | Maximum reward value for species 3 when drawn from a uniform distribution | 0-5 | 0.46 |
| species-4-max | Maximum reward value for species 4 when drawn from a uniform distribution | 0-5 | 0.44 |
| species-1-min | Minimum reward value for species 1 when drawn from a uniform distribution | 0-5 | 0.4 |
| species-2-min | Minimum reward value for species 2 when drawn from a uniform distribution | 0-5 | 0.38 |
| species-3-min | Minimum reward value for species 3 when drawn from a uniform distribution | 0-5 | 0.36 |

| | | | |
|---|---|---|---|
| species-4-min | Minimum reward value for species 4 when drawn from a uniform distribution | 0-5 | 0.34 |
| nectar-max-1 | Maximum reward plants of species 1 can hold | 0-5 | 5 |
| nectar-max-2 | Maximum reward plants of species 2 can hold | 0-5 | 5 |
| nectar-max-3 | Maximum reward plants of species 3 can hold | 0-5 | 5 |
| nectar-max-4 | Maximum reward plants of species 4 can hold | 0-5 | 5 |
| percent-pesticide | Proportion of plants with pesticide | 0-1 | 0 |
| seed-success | Percentage of seeds that become plants next season | 0-1 | 0.1 |
| start-nectar-content | Initial reward content of plants at the start of a season | 0-5 | 0 |
| start-num-species-1 | Initial number of plants of species 1 at the start of a simulation | 0-20000 | 1620 |
| start-num-species-2 | Initial number of plants of species 2 at the start of a simulation | 0-20000 | 1620 |
| start-num-species-3 | Initial number of plants of species 3 at the start of a simulation | 0-20000 | 1620 |
| start-num-species-4 | Initial number of plants of species 4 at the start of a simulation | 0-20000 | 1620 |
| plant-grid? | Setup plant in a uniform grid. This disregards starting seed numbers | TRUE / FALSE | FALSE |
| plant-limit | Maximum number of plants allowed in the field | 0-25921 | 25921 |
| plant-mortality-rate | Proportion of plants from each species that die at the end of every season | 0-1 | 0.2 |
| plant-seeds | Number of potential seeds every plant begins a season with | 0-100 | 6 |

*Table 4: Colony Variables*

| **Variable Name** | **Description** | **Range** | **Default Value** |
|---|---|---|---|
| bee-pop | Number of bees in the colony | 0-200 | 100 |
| colony-nectar-content | Amount of nectar reward stored in the colony | 0-∞ | 0 |

Table 5: General Variables

| General Variables | Description | Range | Default Value |
|---|---|---|---|
| plot-things? | Plot graphs? | TRUE / FALSE | TRUE |
| trace? | Draw the paths of each bee? | TRUE / FALSE | FALSE |
| season-length | Number of ticks in a season | 0-100000 | 6600 |
| stop-on-season? | The simulation pauses at the end of the current season | TRUE / FALSE | FALSE |

## 3. Process Overview & Scheduling

### A. Conceptual Overview

A conceptual overview of SimBee is shown in Figure 4. At the beginning of the first season, individual *bees* are placed randomly in the environment and begin foraging on *plants* of each species (sampling mode). One *plant* per species (if present) is visited per field-of-view. *Bees* store *plant* species and nectar reward from each visited *plant* in working memory as individuals move from *plant* to *plant*. Once working memory capacity is reached, each new species-reward association eliminates the oldest association. *Plants* can vary in their reward values and consequently, *bees* must use reward information stored in working memory to accurately assess the overall reward quality of available species. After completion of sampling mode, *bees* use information present in working memory to decide which *plants* in their field of view to visit in order to maximize reward delivery to the *colony* (foraging mode). Unlike the sampling mode, only one *plant* is visited within each field of view. Since *bees* continuously update reward information in working memory as they visit *plants*, they have the flexibility to change foraging choices if *plant* availability or average reward changes. *Bees* also periodically enter sampling mode in order to acquire reward information on species not visited during the foraging period.

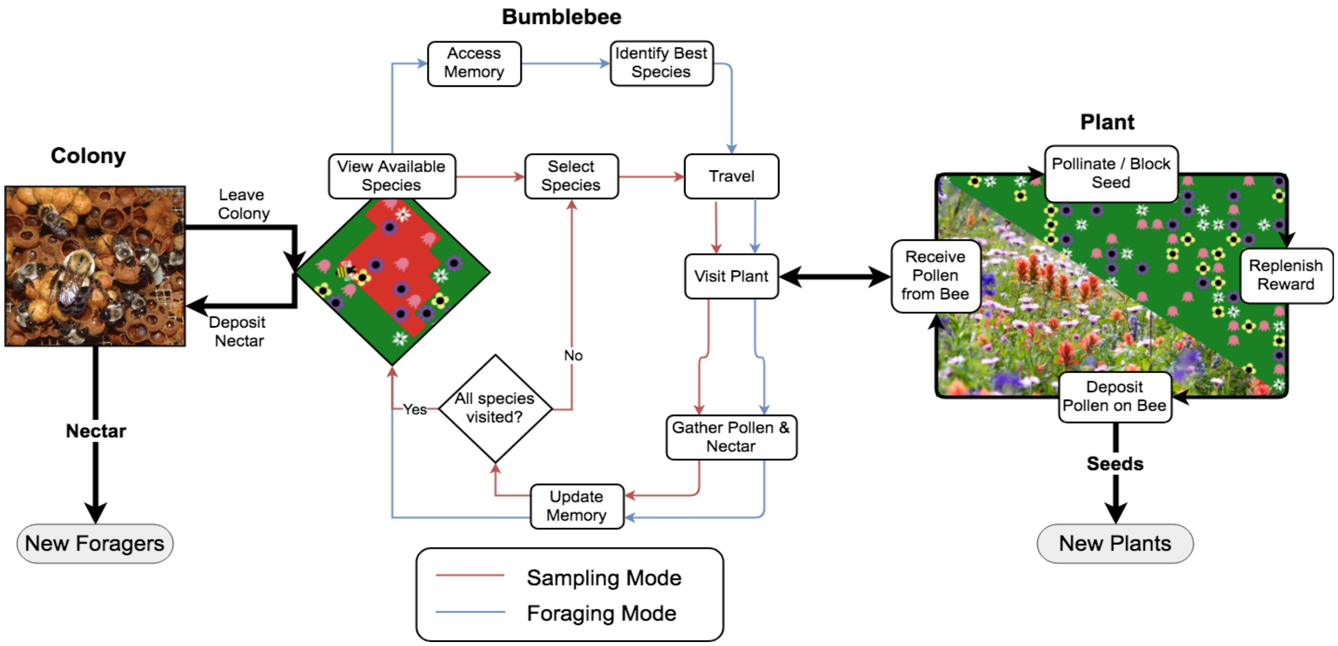

*Figure 4:* ***The SimBee system.*** *Bumblebee foragers leave the colony and enter Sampling mode for approximately 10 floral visits (red lines). The red shaded area in the 'View Available Species' diamond shows the visual range or 'field-of-view' of an individual forager. Foragers sample flowers of different species from the field-of-view, and store reward and species information in working memory for 100 ticks. Foragers then enter Foraging mode (blue lines), where they visit flowers in the field-of-view with the highest reward based on working memory values, which are updated with each visit. After 2000 ticks, they return to Sampling mode, and the cycle continues to the end of the season. Each flower visit made by bees has a reproductive consequence for the plant (seed/no seed produced) that is based on their foraging history (pollen collected from conspecific and heterospecific plants). At the end of the season, colony reward stores and seed totals are used to determine bee and plant abundances at the beginning of the next season, respectively.*

In addition to collecting nectar rewards, *bees* also pick up and deposit pollen as they move from *plant* to *plant*. Thus, visits made by foraging *bees* have reproductive consequences for *plants*. Deposition of a conspecific pollen grain results in fertilization and the production of a seed while deposition of a heterospecific pollen grain prevents the *plant* from receiving a conspecific grain (stigma blockage) and therefore prevents the production of a seed. The likelihood of a pollen grain being deposited on a *plant* is determined in a probabilistic manner following an exponential decay function (detailed below). Lastly, *bees* return to the colony when full of nectar and deposit the nectar and pollen they collected.

### B. Timestep Update

In each time step, each individual *bee* updates its status and performs an action. Typically, the *bee* moves towards a *plant* or collects the reward from a *plant*. This is described in further detail in the bee sub-models section.

*Plants* update only when visited by a *bee*. At the beginning of an interaction with a *bee*, the *plant* updates its reward and pollination status. The *plant* continues to update its reward status each tick while the interaction takes place.

The *colony* updates only when visited by a *bee*. It takes one tick to transfer the collected reward from the *bee* to the *colony*.

### C. Seasonal Update

At the end of the season, unless the bee population size is being held constant, the amount of reward contained in the colony is converted into new foragers that comprise the population the following year. Unless the number of plants is being held constant, the number of new *plants* in the subsequent season is calculated based on the number of seeds produced for each *plant*.

## 4. Design Concepts:

### A. Basic Principles

This agent-based model consists of two primary agents, the *bee* and the *plant*. Each *bee* agent represents an individual and independent bumblebee forager. Each *plant* agent represents a *plant* with a single flower. Since each plant has only one flower with the potential to produce 6 seeds, the terms ‘flower’, and ‘plant’ are interchangeable.

Foraging decisions by *bee* agents seek to maximize nectar reward intake using the agent’s memory of food sources. The decisions of each *bee* agent have a unique impact on the reproductive success of the *plant* agents it visits. Meanwhile, nectar rewards provided by *plant* agents influence foraging decisions by *bee* agents. We hypothesize that exposure to sub-lethal amounts of pesticides affects either the memory capacity or the processing speed of *bee* agents, which in turn will affect bee behavior. The model allows us to test whether impacts of pesticides on the behavior of individual *bee* agents has long term effects on *bee* and *plant* populations over time.

### B. Emergence

Patterns can emerge at both the individual and the population level. At the individual level, patterns emerge in the activities of bees and their foraging decisions; for instance, floral constancy tends to emerge in individual bees even when all flower species are equally rewarding on average, due to memory capacity limitations and reward variability. At the population level, patterns emerge in the growth and decline of *bee* and *plant* populations based on foraging and pollination success. Different patterns emerge when *bee* agent memory is perturbed. For instance, feedback between changes in bee and *plant* populations over multiple seasons can lead to changing patterns of *plant* diversity.

## C. Adaptation

*Bee* agents remember the reward and species of past visits to *plants*. They use these memories to determine which species of *plant* agents to visit at their current location. In addition, bees periodically sample their environment, which allows them to gather information on plant species they may not have visited recently. Thus, individual *bees* can adapt their decisions to a changing environment.

## D. Objectives

Foraging decisions by *bee* agents seek to maximize nectar reward intake using the agents' memory of food sources.

Sampling by *bee* agents aims to improve the accuracy of the bee's memory of the environment by making sure all plant species' rewards are assessed periodically.

## E. Learning

Foraging *bee* agents remember the rewards of past visits and average these to determine the most rewarding species of *plant*. The greater the memory capacity of a *bee*, the more visits it will remember and the more accurate the averages will be. Therefore, the more memory a *bee* has, the better its prediction of the local environment will be, assuming a stable environment within a season.

## F. Sensing

*Bees* perceive the surrounding environment in a radius of 5 patches and an arc of 117 degrees in front of them that we refer to as the *field-of-view*. They can identify the species of *plants* and whether or not they are occupied by other *bees*. Inherent in this implementation is the assumption that *bees* can always identify *plant* agents regardless of any *bee* impairments. *Bees* also perceive time and require a set number of *ticks* to extract rewards from *plants*; this can be adjusted to represent impairments to processing speed.

## G. Interaction

*Bees* interact with *colonies* only to deposit nectar reward and must be located on the same patch to interact. *Bees* and *plants* interact directly. When a *bee* travels to a *plant*, the *bee* can 'land' on the *plant* and extract a nectar reward from it. This interaction also results in the conspecific or heterospecific transfer of pollen. *Bees* also interact indirectly with other *bees* by blocking them from visiting the same *plant*.

## H. Stochasticity

Generally, agents are processed in a random order. The following processes also contain some degree of stochasticity, as described in section 7 submodels:

- *Plant* location
- *Plant* mortality
- *Plant* reward quantity
- *Plant* seed success
- *Bee* starting location
- *Bee* movement orientation
- *Bee* foraging
- *Bee* sampling

### I. Collectives

*Bees* belong to a certain population that deposits accumulated nectar rewards at a single *colony*. Every *bee* in a population is a forager, since we do not model reproductives. The *colony* uses the cumulative reward over a season to produce new *bees* for the next season.

*Plants* belong to a species that shares the same color, shape, and nectar reward distribution. *Plants* contain potential seeds and produce species-specific pollen. Conspecific pollen transfer results in production of a seed that can produce a *plant* of the same species.

### J. Observation

The interface shows a map of the modelled landscape, as well as the location of *plants*, *bees*, and *colonies*. Plots provide a number of outputs such as population size, numbers of seeds produced, average nectar reward, nectar collected, etc. Additional output can be recorded through the built-in Behaviorspace.

## 5. Initialization

*SimBee* is initialized by the function *setup*. This function clears all agents and global variables, calls additional functions to initialize the environment and agents, and resets the time step to 0. Next, the *colony* is initialized and placed in the environment by a function called *setup-colonies* at the center of the grid (coordinate 0,0). Afterwards, *bees* are created by the function *setup-bees*. This function creates *bees* for each *colony* in the environment and initializes *bee*-specific variables. *Plants* are initialized after *bees* using the *setup-plants-random* function. *Setup-plants-*

*random* places *plants* in random empty patches following a random uniform distribution. Lastly, the function *setup-seasonal-globals* initializes season-specific global variables used for population-level statistics and visualization.

A simulation starts at the beginning of the spring season with an initial number of bumblebee foragers and seeds. The appropriate number of patches are chosen at random on the environmental grid and filled with one *plant* each. *Bees* are also randomly distributed around the environmental grid and set to sampling mode. Default values of all state variables are shown in the tables above.

## 6. Input Data

The model does not include any input data.

## 7. Submodels

### A. Model Parameters

Model parameters such as handling time, nectar reward amounts, flower handling times, etc. are based on experimentally observed data for species of bumblebees and *plants* commonly found in the Eastern United States. The user can change many of these values via sliders to determine realistic values for other desired scenarios. Refer to Table S5 for the model variables used for each scenario in the associated publication.

*Table 6: Model variable values for each Scenario Observer (output) variables*

| **General Variables** | **Description** | **Scenario 1** | **Scenario 2** | **Scenario 3** |
|---|---|---|---|---|
| season-length | Number of ticks in a season | 6600 | 6600 | 6600 |

|  |  |  |  |  |
|---|---|---|---|---|
| **Colony Variables** | **Description** | **Scenario 1** | **Scenario 2** | **Scenario 3** |
| number-of-colonies | Number of bee colonies | 1 | 1 | 1 |
| bee-nectar-max[2] | Maximum units of reward a bee can hold before needing to return to the colony | 80 | 80 | 80 |
| max-rand-visit-chance | Probability of a bee with zero memory visiting a plant when available | 1 | 1 | 1 |
|  |  |  |  |  |
| **Bee Variables** | **Description** | **Scenario 1** | **Scenario 2** | **Scenario 3** |
| constant-bees? | Population of each hive remains constant between seasons | FALSE | TRUE | FALSE |
| nectar-per-reproductive[3] | Amount of nectar required to produce a new bee for the next seasons | 402 | 329 | 329 |
| colony-1-add-time-on-plant-for-impaired[4] | Additional number of ticks a bee from colony 1 impaired by pesticides must spend on plants | 0,1,2 | 0,1,2 | 0,1,2 |
| colony-1-impaired-memory-size | Number of past visits a bee from colony 1 impaired by pesticides can remember | 0 | 0 | 0 |
| colony-1-normal-memory-size | Number of past visits a bee from colony 1 can remember | 10 | 10 | 10 |
| colony-1-number-of-bees | Number of bees in colony 1 | 100 | 100 | 100 |
| colony-1-percent-susceptible | Proportion of bees in colony 1 susceptible to pesticides | 0, 0.25, 0.5, 0.75, 1 | 0, 0.25, 0.5, 0.75, 1 | 0, 0.25, 0.5, 0.75, 1 |
| colony-1-close-enough | On a scale from 0 to 1, what percent difference in reward can bees from colony 1 distinguish? | 0.07 | 0.04 | 0.04 |
|  |  |  |  |  |
| **Sampling Variables** | **Description** | **Scenario 1** | **Scenario 2** | **Scenario 3** |
| intra-sample-duration[5] | Number of ticks all bees sample for during a sampling bout | 100 | 100 | 100 |
| use-uniform-random-sampling | Should the time of the second sampling bout of all bees be drawn from a uniform random distribution? | FALSE | FALSE | FALSE |

---

[2] Unpublished data in figure S5 informed our choice of 80 units of nectar per foraging trip for a single bee.

[3] Data collected indicate that there is a linear correlation between resource intake and number of reproductive individuals produced[131].

[4] Immune impairments to bumblebees can cause a 40% reduction in foraging efficiency through reduced cognitive flexibility and reward intake rates[39].

[5] Heinrich demonstrates that bumblebees periodically sample their environment to optimize their foraging behavior[79,132].

| | | | | |
|---|---|---|---|---|
| resample-deviation | The standard deviation of values drawn for the number of ticks between sampling bouts | 0 | 0 | 0 |
| resample-rate[6] | The mean number of ticks between sampling bouts | 2000 | 2000 | 2000 |
| sample-duration-deviation | The standard deviation of values drawn for the duration of sampling bouts | 0 | 0 | 0 |
| | | | | |
| **Plant Variables** | **Description** | **Scenario 1** | **Scenario 2** | **Scenario 3** |
| constant-plants? | Population of each species of plant remains constant between seasons | TRUE | FALSE | TRUE |
| instant-refill?[7] | Instantly regenerate reward in plants when visited by a bee | TRUE | TRUE | TRUE |
| use-normal-distribution? | If *instant-refill?* is TRUE, should plant rewards refill based on a standard normal distribution (TRUE) or uniform random distribution (FALSE)? | FALSE | FALSE | FALSE |
| max-1[8] | Maximum reward value for species 1 when drawn from a uniform distribution | 0.95 | 0.68 | 0.68 |
| max-2 | Maximum reward value for species 2 when drawn from a uniform distribution | 0.85 | 0.68 | 0.68 |
| max-3 | Maximum reward value for species 3 when drawn from a uniform distribution | 0.5 | 0.68 | 0.68 |
| max-4 | Maximum reward value for species 4 when drawn from a uniform distribution | 1 | 0.68 | 0.68 |
| min-1 | Minimum reward value for species 1 when drawn from a uniform distribution | 0.45 | 0.48 | 0.48 |
| min-2 | Minimum reward value for species 2 when drawn from a uniform distribution | 0.35 | 0.48 | 0.48 |
| min-3 | Minimum reward value for species 3 when drawn from a uniform distribution | 0.5 | 0.48 | 0.48 |
| min-4 | Minimum reward value for species 4 when drawn from a uniform distribution | 0 | 0.48 | 0.48 |
| percent-pesticide | Proportion of plants with pesticide | 1 | 1 | 1 |
| seed-success[9] | Percentage of seeds that become plants next season | NA | 0.4 | 0.4 |
| start-num-species-1 | Initial number of plants of species 1 at the start of a simulation | 4000 | 1000 | 1000 |

---

[6] The resample-rate was chosen based on the season length such that bees resample their environment three times during the virtual season. This was based on Heinrich 1979[79] and Mobley & Gegear 2018[39].

[7] Bee-adapted species of plants quickly replenish a small amount of concentrated nectar[133].

[8] Plants vary in reward and secretion rate. The values we use for min and max reward fall into the ranges suggested by Castellanos et al. 2002[133].

[9] We use the lower-bound of 40% survival from Metz et al. 2010[134].

| | | | | |
|---|---|---|---|---|
| start-num-species-2 | Initial number of plants of species 2 at the start of a simulation | 4000 | 1000 | 1000 |
| start-num-species-3 | Initial number of plants of species 3 at the start of a simulation | 4000 | 1000 | 1000 |
| start-num-species-4 | Initial number of plants of species 4 at the start of a simulation | 4000 | 1000 | 1000 |
| plant-limit[10] | Maximum number of plants allowed in the field | 25900 | 25900 | 25900 |
| plant-mortality-rate[11] | Proportion of plants from each species that die at the end of every season | 0 | 1 | 1 |
| plant-seeds[12] | Number of seeds every plant begins a season with | NA | 6 | 6 |
| | | | | |
| **Pollen Variables** | **Description** | **Scenario 1** | **Scenario 2** | **Scenario 3** |
| pollen-model[13] | Specifies the model for pollen fall-off / decay once it is collected by bees | exponential | exponential | exponential |
| pollen-grains-transferred[14,15] | Number of pollen grains a bee picks up from a plant during a single visit | 40 | 40 | 40 |
| exp-decay-rate[16] | Value for the exponential decay model | 0.7 | 0.7 | 0.7 |

## B. Colony Agent

As with real bumblebees, *colony* agents store all floral rewards collected by *bee* agents. *Bee* agents spend one time-unit in the *colony* between foraging bouts and deliver rewards determined by *bee-nectar-max*. At the end of a season, the colony agent converts its total reward content into new *bee* agents for the next season with a conversion rate set by *nectar-per-reproductive*. The *colony* does not use any reward during the season.

---

[10] The plant-limit is based on the spatial limitations of the virtual environment and is not reflective of a specific foraging environment.

[11] In our simulations plants with a 0% mortality rate are considered perennial while plants with a 100% mortality rate represent annuals.

[12] The number of seeds per flower can vary dramatically by species. We believe six seeds is reasonable[131].

[13] Pollen model identified in Rademaker et al. 1997[68].

[14] Pollen pick-up informed by Cane JH, Dunne R. 2014[135] and Castellanos et al. 2003[136].

[15] Total pollen produced by plants varies by species and was informed by Mondal & Mandal 1998[137].

[16] Pollen deposition and loss informed by Cane & Dunne 2014[135].

## C. Bee Submodels

### *I. Field of view*

Individuals have a visual range for *plant* detection that has a radius of 5 patches and an arc of 117 degrees (see Figure 4). This corresponds to a view of approximately 25 patches or 0.096% of the simulated environment. Thus, individuals have the potential to detect and make decisions on a maximum of 25 *plants* at a time.

### *II. Memory*

All individuals have both working memory (information storage and processing from seconds to minutes timeframes) and long-term memory (information storage from days to weeks). We assume that individuals remember that they obtained reward from all four *plant* species in the past as well as the location of the colony and their relative position to it at all times (long-term memory). In our model, we store this long-term information as individual parameters that do not change over the life of the bee. We also assume that individuals use working memory while foraging to store reward information associated with each flower visit and to decide on which *plant* species to visit to maximize reward delivery to the colony. As it relates to foraging bees, we define working memory as the ability to remember reward values associated with flowers visited in the recent past (seconds to minutes). Working memory is structured as a list of reward-plant species pairs, with 20 being the maximum number of pairs in the list (*max-memory-length*). In the model, this is implemented as a list of tuples called *visited-memory*. A memory capacity of 10 was used for the scenarios in the associated work. Each individual starts a season with an empty memory list and continues to add reward-species pairs with each visit. Once the working memory capacity is reached, each new reward-species pair added to the list deletes the oldest memory from the list. Floral reward information in working memory is maintained between foraging bouts. An individual *bee* agent with a working memory capacity of 0 (no memory) randomly selects one flower in its *field-of-view* to visit when in foraging mode.

### *III. Plant handling time*

Handling time is the time individual *bee* agents spend in contact with the *plant* while collecting a reward. *Plant* handling times range from 5 to 7 time-units. For all simulations the baseline *plant* handling time, regardless of reward level, is 5 units. Additional time spent handling *plants* is used to simulate stressor-induced cognitive impairments in information processing speed.

### *IV. Movement*

The *bee* agent follows several movement rules. The *bee* will always have one of three destinations: a *plant*, a *colony*, or a new *field-of-view*. The agent moves directly towards its destination moves directly to its destination, either by moving the side-length of one patch per tick, or, if the center of the destination patch is less than one patch-length away, to that center. If the agent cannot find a *flower* for its destination, it rotates randomly within a 60º arc and sets its destination to be the furthest patch at the edge of its field of view. Additionally, if an agent reaches the edge of the simulated environment, it will randomly turn until it is no longer facing the edge.

When a *bee* agent is full of nectar and sets its destination to be the *colony*, the agent takes one tick to teleport to the *colony* rather than traversing the distance at one patch per tick. After depositing the nectar (an action that takes one tick), the *bee* agent takes one tick to teleport to its previous position in the environment. These movements are performed by ‘teleporting’ so that essentially no time is spent traveling to and from the colony. We implemented this to remove the effect of long travel time due to starting location.

### *V. Sampling mode*

In order to acquire and update information on the reward properties of available *plant* species, individuals periodically (at the beginning of each season and every 2,000 ticks thereafter) sample *plants* of each *plant* species in its *field-of-view* for 100 ticks. Such sampling behavior enables foragers to detect any changes in the quality and availability of *plant* resources, which is a key component of adaptive foraging under variable resource conditions[138]. While sampling *plants*, a *bee* will visit one of each species of *plant* in its *field-of-view*, chosen at random, before creating a new field of view and repeating the process (Figure 5). If there are no *plants* in its *field-of-view,* then the *bee* will travel to the edge of its *field-of-view* and search again. Consecutive *field-of-views* are also spatially separated from one another so that available flowers in consecutive *field-of-views* are different. See Table 7 for sampling-specific variables and Table 6 for the values used in our experiments.

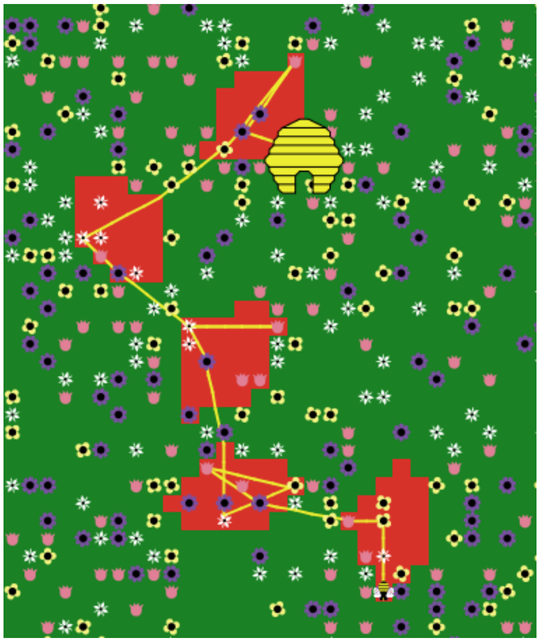

*Figure 5: Example of an individual bee in sampling mode after leaving the colony. The red patches represent each consecutive field-of-view of the bee.*

*Table 7: Sampling Variables*

| **Sampling Variables** | **Description** | **Range** | **Default** |
|---|---|---|---|
| usesampling? | Should all bees regularly sample the environment? | TRUE / FALSE | TRUE |
| intra-sample-duration | Number of ticks all bees sample for during a sampling bout | 0-2000 | 100 |
| num-plants-to-sample | Number of plants of each species to sample within a single field of view | 0-10 | 1 |
| use-uniform-random-sampling | Should the time of the second sampling bout of all bees be drawn from a uniform random distribution? | TRUE / FALSE | TRUE |
| resample-deviation | The standard deviation of values drawn for the number of ticks between sampling bouts | 0-100 | 0 |
| resample-rate | The mean number of ticks between sampling bouts | 0-50000 | 1000 |
| sample-duration-deviation | The standard deviation of values drawn for the duration of sampling bouts | 0-100 | 0 |

### *VI. Foraging mode*

According to optimal foraging theory, individuals should make foraging decisions that maximize their rate of reward intake given information[139,140]. Bumblebees base their decisions on a combination of reward, handling time, and traveling time. In our model, we chose to use a simpler decision algorithm that relies only on a *bee's* memory of *plant* rewards to determine the probability of a visit. The bumblebee begins foraging mode by determining which *plant* species in its *field-of-view* to visit based on its memory of the environment (Figure 4). It then chooses a *plant* of that species based on a decision algorithm (*Decision-Making*) and travels to it (Figure 6). Once the *bee* reaches the *plant*, it lands on the *plant* and extracts a nectar reward from it; this process requires several ticks. The *bee* also transfers pollen (*Pollen Pickup and Loss, Pollen Deposition*). The *bee* then updates its memory list by adding the nectar reward value and *plant* species pair corresponding to this visit and repeats the process using a new *field-of-view* from its current location on the *plant*.

In some cases, the *bee* may choose not to visit any *plants* in its *field-of-view* (see *Decision-Making*). In that case, the *bee* travels to the edge of its current *field-of-view* and creates a new *field-of-view* at that location.

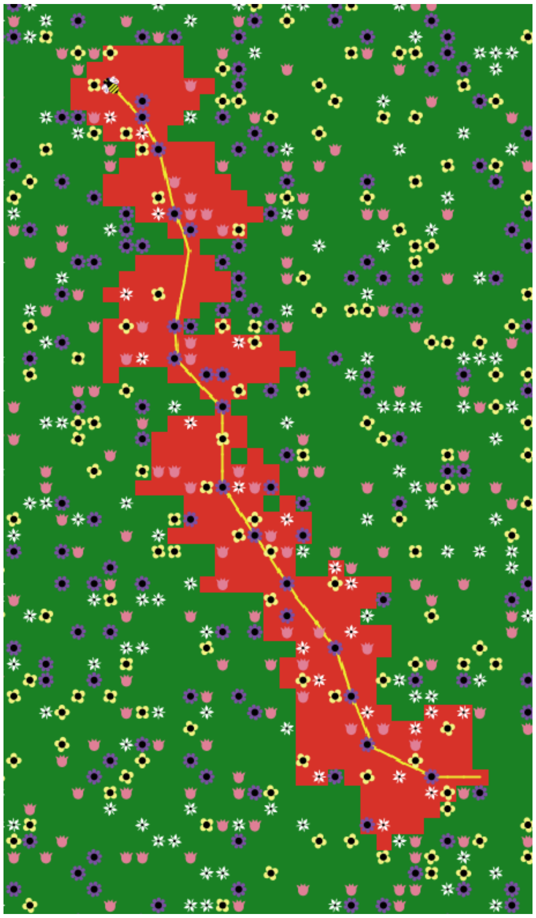

*Figure 6: Example of an individual bee in foraging mode. The red patches represent each consecutive field-of-view of the bee. Each new field-of-view is generated from the plant that the bee visits in its current field of view, or from the edge of the previous field-of-view if no plant is chosen to visit.*

### *VII. Choose-Plant*

According to optimal foraging theory, individuals should adopt plant choice patterns that maximize their rate of reward intake[139,140]. In highly variable reward environments, this ability depends on the amount of past reward information stored in memory - more individual values yield a more accurate estimate of average floral rewards at the species level. SimBee therefore uses a memory-based algorithm to determine the probability of visiting all plant species in view.

Foragers use a memory-based algorithm to determine the probability of visiting each plant species for a given *field-of-view*. The plant species determined to have the greatest reward are always taken if present. If not, then the forager will visit the plant species in the *field-of-view* offering the highest average reward, with a probability based on its reward relative to the 'best' species. To avoid placing too much weight on outlier rewards, median rather than mean reward values in memory are used. Each step of the algorithm is detailed below.

**Step 1: Calculate median reward for each plant species in working memory.**

Let the length of memory be *n* and the number of plant species be *S*. Let each visit to a species be $s_i$ where $i = 1,2, \dots, S$ and the corresponding reward be $r_k$ where $k = 1$ (most recent value in memory), $2, \dots, n$ (oldest value in memory). Let the median reward value for each species be $\tilde{r}_i$ for $i = 1, \dots, S$.

**Step 2: Determine probability of visitation.**

Let $y_i$ be the scaled median reward:

$$y_i = \frac{\tilde{r}_i}{\sum_1^S \tilde{r}_i} \text{ for } i = 1, \dots, S$$

*Equation 1*

Let *Probability of visitation* ($Pa_i$) be defined as

$$Pa_i = \frac{y_i}{\max y_i} \text{ for } i = 1, \dots, S$$

*Equation 2*

If the bee has no memory of a species, reward is treated as zero and the ensuing probability of visiting that species will also be zero.

**Step 3: Determine plant species with the highest reward quality.**

The plant species with the highest $Pa_i$ is placed on the list $Best_a$. If the $Pa_i$ of another plant species falls within a predetermined probability range called *Close Enough* (*ce*) then that species is also added to $Best_a$ (i.e. the normalized values are considered to be the same).

**Step 4: Select a plant**

All plant species in the field-of-view are stored in a list called *Species in View* $(Sv)$. The probability of visiting plant species in $Sv$ is determined by

$$Pv_i = \frac{y_i}{\max y_i} for\ all\ i\ in\ Sv$$

*Equation 3*

As with $Pa_i$, if the bee has no memory of the species $i$ then $Pv_i = 0$. The plant species with the highest $Pv_i$ and the species with $Pv_i > P\ v_i - ce$ are then stored in a list called $Best_v$. If a plant species in $Best_v$ is also in $Best_a$ (meaning it is one of the best species available based on current values in working memory), then one plant of that plant species is randomly chosen from the field-of-view. If a plant species in $Best_v$ is not in $Best_a$, then one flower of that species will be chosen with the probability $Pa_i$.

Note that if all species in the field-of-view are view are not currently in $Best_v$ and $Best_a$, then the probability of visitation is zero. That is, if a species is not in a bee's working memory, the bee will not visit that species again until it next enters sampling mode.

If no plant species is chosen to visit, then the bee chooses a destination at the edge of its current *field-of-view* and moves toward that destination rather than toward a chosen plant.

After the bee reaches its chosen destination, a new *field-of-view* is determined. Steps 1-4 are repeated for all subsequent field-of-views in foraging mode.

### *VIII. Reproduction*

At the end of the season, the amount of reward contained in the colony is converted into new foragers that comprise the population the following year. See Table 2 for the reward-to-bee conversion factors used in different experiments. We assume that all energetic demands of individual foragers are being met; thus, all floral rewards contained in the colony are used to determine its reproductive output.

## D. Plant Sub-models

### *I. Plant Distribution*

At the beginning of each season, *plants* from up to four plant species are randomly placed on empty patches on the environmental grid. For seasons 2-20, if the sum of seed totals for all plant species is greater than the number of empty patches, then the model will discard all excess seeds in proportion to the number generated for each species.

### *II. Plant Reward*

Reward distributions of all *plant* species follow a uniform distribution bounded by a specified minimum and maximum value. Rewards are completely drained during a *bee* visit and instantly refilled with a value randomly drawn from the distribution upon *bee* arrival; thus, each *plant* has a floral reward that changes around the fixed average reward for that species with each *bee* visit.

*III. Pollen Transfer*

Bees pick up 40 pollen grains per floral visit and deposit all pollen in the colony every time they return from a foraging run (pollen levels on the forager are reset to 0). During a foraging run, bees deposit pollen on flowers as a function of the pollen load they are carrying. The dynamics of pollen deposition on a sequence of flowers follows the exponential decay model reported in previous field studies[68].

The pollinator loses pollen grains it acquired previously as a function of the pollen it is carrying. Pollen load is re-calculated independently for each species after each floral visit. If we let $v_j$ be the number of visits since the pollinator last visited *plant* species $j$, and $l_{i,j}$ is the pollen load from *plant* species $j$ on the pollinator for visit $i$, then the pollen load for plant species $j$ for visit $i + 1$ is calculated as:

$$l_{i+1,j} = l_{i,j} * e^{\{-r * v_j\}}$$

*Equation 4*

where $0 < r < 1$ is a constant value. For our experiments, we chose $r = 0.7$ such that the pollen load from one visit to a specific species will approach zero after five visits to other species. Thus, after five visits, total pollen grains on a pollinator reaches an equilibrium value. See Table 8 for a detailed list of the variables used. Refer to Table 6 for the values used in our experiments.

*Table 8: Pollen Variables*

| **Pollen Variables** | **Description** | **Range** | **Default** |
|---|---|---|---|
| pollen-model | Specifies the model for pollen fall-off / decay once it is collected by bees | Constant, Linear, Exponential | exponential |
| pollen-grains-transferred | Number of pollen grains a bee picks up from a plant during a single visit | 0-100 | 40 |
| exp-decay-rate | Value for the exponential decay | 0-1 | 0.7 |

| | | | |
|---|---|---|---|
| | model | | |

*IV. Reproduction*

Each plant has the potential to produce six seeds. Deposition of conspecific pollen results in the production of a seed (pollination), while deposition of heterospecific pollen results in the elimination of an empty seed slot (stigma blockage). The probability of a pollination event is equal to the proportion of the conspecific pollen carried by the bee when it lands on the flower.

$$Probability\ of\ pollination = l_i / \sum_{k=1}^{S} l_k$$

*Equation 5*

When all six seed slots of a plant have been used up, pollen deposition and pickup processes continue without seed production.

*V. Determination of population size*

At the end of the season, unless the plant population is being held constant, 40% of the total number of seeds produced are randomly selected and placed in the environment to start the next season (a 60% seed mortality rate). Plants are randomly re-positioned at the beginning of each season to control for any effects of plant distribution on bee foraging patterns. The size of the environment limits the max number of plants (flowers) in the system to 25921 individuals.

## 8. Supporting Data

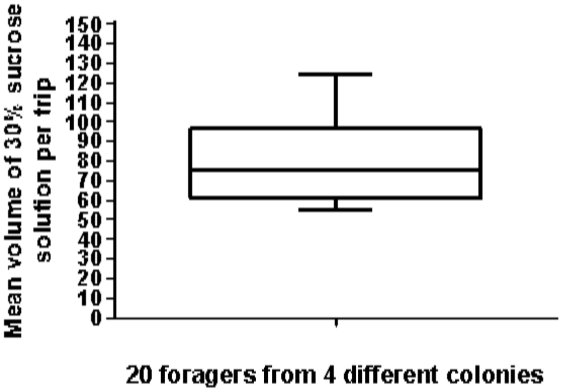

*Figure 7: Unpublished results from lab experiments show that bumblebee Bombus impatiens foragers return to the hive with approximately 80 mg of reward per trip.*

### Methods

*Experimental Setup: Bombus impatiens* colonies were obtained from Biobest Biological Systems Canada (Leamington, ON) and connected to a small flight cage (325 X 240 X 221cm) with a gated tube constructed from wire mesh, thus enabling control over the number of bees entering the cage. The cage was illuminated by two Ultra SunTM 6500K (ZooMed Laboratories Inc., San Luis Obispo, USA) and two Sylvania GRO-LUX fluorescent lights. Prior to experiments, foragers collected 30% sucrose solution from several feeders placed in the center of the cage. All foragers within a colony were marked for identification with different color combinations of acrylic paint. Colonies were directly supplied with pollen ad libitum to facilitate nectar foraging during experiments.

*Flowers and test array:* Artificial flower types (herein referred to as 'flowers') were constructed by removing the cap from clear 1.5mL Eppendorf centrifuge tube and fixing a 3cm (diameter) circular collar made of white CreatologyTM foam (Michaels Stores, Inc.) around the entrance of the tube. Flowers were presented to bees in an array consisted of a horizontal 120 X

80 X 5cm Styrofoam block covered a color print of natural foliage. At total of 96 holes were drilled in the block to hold artificial flowers, which were distributed 12 rows of 8 (12 cm apart within rows and 6 cm between rows) with adjacent rows offset by half the distance between flowers in each row. Each flower contained 2 uL of 30% sucrose solution deposited at the base of the tube.

*Data collection and analysis:* Marked foragers were digitally recorded for two foraging trips on the test array. Only one forager was tested at a time. Flowers visited by foragers were checked to ensure that they had been drained and then immediately refilled with 2 uL of 30% sucrose solution, which occurred once the forager had entered the tube of the subsequent flower. A total of 20 bees from 4 colonies (5 bees per colony) were tested in this way. The volume of test solution collected per foraging trip (number of flowers visited X volume of solution per flower) was determined from the digital recordings of each forager. The volumes of the two foraging trips were averaged to yield a single value for each forager.

# Chapter 3: Modeling scale up of anthropogenic impacts from individual pollinator behavior to pollination systems

The work presented in this chapter was published in Volume 35, Issue 5 of Conservation Biology and is partially reproduced below.

Gegear, R. J[*]., Heath, K. N[*]. & Ryder, E. F. Modeling scale up of anthropogenic impacts from individual pollinator behavior to pollination systems. *Conservation Biology* **35**, 1519–1529 (2021).

The research that led to this publication was a collaborative effort between members of the Ryder lab at Worcester Polytechnic Institute and the Gegear lab at UMass Dartmouth. Robert Gegear provided expertise in ecology and bumblebee research as well as designing research questions. Kevin Heath designed, documented, and validated the SimBee model, and generated data for all scenarios considered in the manuscript. Elizabeth Ryder also designed research questions and edited drafts of the manuscript. All authors analyzed data generated and reviewed and approved the final version of the manuscript.

## Abstract

Understanding how anthropogenic disturbances affect plant–pollinator systems has important implications for the conservation of biodiversity and ecosystem functioning. Previous laboratory studies show that pesticides and pathogens, which have been implicated in the rapid global

---

[*] These authors contributed equally to this work

decline of pollinators over recent years, can impair behavioral processes needed for pollinators to adaptively exploit floral resources and effectively transfer pollen among plants. However, the potential for these sublethal stressor effects on pollinator-plant interactions at the individual level to scale up into changes to the dynamics of wild plant and pollinator populations at the system level remains unclear. We developed an empirically parameterized agent-based model of a bumblebee pollination system called SimBee to test for effects of stressor-induced decreases in the memory capacity and information processing speed of individual foragers on bee abundance (scenario 1), plant diversity (scenario 2), and bee–plant system stability (scenario 3) over 20 virtual seasons. Modeling of a simple pollination network of a bumblebee and four co-flowering bee-pollinated plant species indicated that bee decline and plant species extinction events could occur when only 25% of the forager population showed cognitive impairment. Higher percentages of impairment caused 50% bee loss in just five virtual seasons and system-wide extinction events in less than 20 virtual seasons under some conditions. Plant species extinctions occurred regardless of bee population size, indicating that stressor-induced changes to pollinator behavior alone could drive species loss from plant communities. These findings indicate that sublethal stressor effects on pollinator behavioral mechanisms, although seemingly insignificant at the level of individuals, have the cumulative potential in principle to degrade plant–pollinator species interactions at the system level. Our work highlights the importance of an agent-based modeling approach for the identification and mitigation of anthropogenic impacts on plant–pollinator systems.

## Methods

### Model Overview

We used the NetLogo[130] platform to create SimBee, a virtual bumblebee pollination system. Individual virtual bees store floral reward information about different plant species in working memory. Bees use this information to adaptively and rapidly collect resources for the colony (bee fitness) under highly variable floral resource conditions, transferring pollen among four co-flowering plant species in the process (plant fitness) over multiple virtual seasons. Bee–flower interactions at the individual level thus dictate the population dynamics and diversity of the system. Full details of SimBee are provided in Chapter 2: SimBee ODD, following the ODD protocol[129,141,142], including references for parameterization of the model. See the publication Appendix S3 for the model code.

### Bees

Bees move through the virtual environment visiting and collecting rewards from flowers in consecutive field of views (FOVs) (approximately 0.096% of the simulated foraging environment). For each virtual season, bees begin in sampling mode, where they randomly visit one flower of each plant species present in each FOV. Bees update their memory by storing the floral reward collected and plant species identity for each visit (memory capacity = 10 reward–species pairs). Bees next enter foraging mode and visit what they decide is the most rewarding flower species in each FOV. The probability of a visit is based on the average reward values for each plant species in memory (details of the decision-making algorithm in Chapter 2: SimBee ODD). As bees forage, information is updated in memory. Once capacity is reached, each new value pair replaces the oldest pair. Bees alternate between sampling and foraging modes for the rest of the virtual season (5% sampling mode, 95% foraging mode). Each bee collects 80 units of nectar reward per foraging bout and deposits all rewards (nectar and pollen) in the colony between bouts. Total reward in the colony at the end of the season determines the size of the bee population the following virtual season, unless bee populations are held constant. Bees with impaired memory capacity fail to store information in memory; therefore, they randomly select one flower from each FOV to visit. Bees with decreased information processing speed take 40% longer per flower visit.

### Plants

Reward levels of each virtual plant species follow a uniform distribution bounded by a specified minimum and maximum value. Upon bee arrival at the flower, reward is replenished with an amount randomly drawn from the distribution; thus, each plant has a floral reward that varies around the fixed average reward for that species. Bees pick up 40 grains of pollen per flower visit and deposit them during subsequent flower visits following the exponential decay function in Rademaker et al. (1997)[68]. Each plant has the potential to produce one seed per bee visit up to a maximum of six seeds per plant. Deposition of conspecific pollen results in the production of a seed (pollination), whereas deposition of heterospecific pollen results in the elimination of an empty seed slot (stigma blockage). The probability of a pollination event is equal to the proportion of the conspecific pollen carried by the bee when it lands on the flower. At the end of each virtual season, a random selection of 40% of total seeds produced by plants determines population size for each plant species during the next virtual season, unless plant populations are held constant.

## Model design assumptions

The SimBee system was designed to include only elements essential to address our central question of whether stressor induced changes to forager behavior can scale up from the individual to the system level. Consequently, we did not include known lethal and sublethal stressor effects on other aspects of the bumblebee life cycle, which would only increase deleterious effects on our bee–plant system. We abstracted the critical features of individual bee–plant interactions into key model processes governing forager behavior, memory, decision-making, and pollen transfer. The most important design concept in the model is the ability of foragers to use memory to adapt their behavior to maximize fitness in a varying floral environment. We assumed that bees remember receiving floral rewards from each plant species in the past, but rely on information stored in working memory to obtain and act on current average reward values associated with each species. Although our virtual season does not currently include natural seasonal effects, such as floral blooms, it does allow for testing of different kinds of variation in floral environments that might arise during a real-world season, such as equal or unequal average floral rewards and plant species abundances. We also assumed that sublethal stressor effects influence individual bee–plant interactions over the entire virtual season and between virtual seasons. We believe this is a reasonable assumption because nectar and pollen contaminated with pesticides is stored in the colony for long periods, thereby providing a route of oral exposure to new foragers over the colony cycle[15]; pathogens can be passed between cohorts throughout colony development[143]; acute exposure to pesticides and infection can permanently impair bee cognitive abilities[37,39]; and pesticides and high infection rates can persist in foraging environments for multiple seasons[144–147].

*Table 9: Experimental model scenarios to test stressor effects on bee abundance, plant diversity, or system stability*

| **Scenario** | **Initial no. of bees** | **Initial no. of plants/ species** | **Mean plant rewards (min, max)** | **No. of bees over virtual seasons** | **No. of plants over virtual seasons** |
|---|---|---|---|---|---|
| 1. Bee abundance | 100 | 4000 | Sp 1 0.7(0.45,0.95)<br>Sp 2 0.6 (0.35,0.85)<br>Sp 3 0.5 (0.5,0.5)<br>Sp 4 0.5 (0.0,1.0) | changes based on rewards gathered | constant |
| 2. Plant Diversity | 100 | 1000 | All species<br>0.58 (0.48, 0.68) | Constant | changes based on seeds produced |
| 3. System Stability | 100 | 1000 | All species<br>0.58 (0.48, 0.68) | changes based on rewards gathered | changes based on seeds produced |

## System patterns in the SimBee model

To consider our model realistic enough for its purpose, we confirmed that emergent patterns resulting from individual bee-plant interactions matched those present in natural systems (Appendix S1: Validation of Model Assumptions). Virtual bees used floral reward information stored in memory to make foraging decisions that were adaptive for the colony (Appendix S1: Validation of Model Assumptions). Bee foraging (pollen transfer) patterns had a direct effect on plant reproductive success. In addition, we found that virtual bees with normal memory exhibited floral constancy, just as real bees do[67], making like–like floral moves in more than half of their floral visits, even when all four floral species were equally rewarding. In contrast, bees lacking memory visited floral species at random, with a consequent reduction in seed production (Appendix S1: Validation of Model Assumptions). Thus, plant reproductive success depended as expected on bee foraging behavior.

## Experimental Scenarios

We investigated effects of stressor-induced changes to bee cognition and foraging behavior on bee abundance and plant diversity under three ecologically relevant model scenarios (Table 9). The degree of each type of cognitive or behavioral impairment in foragers was based on previous laboratory work[35–37,39] and included either an inability to remember the reward properties of available flowers (decreased memory capacity) or a 40% delay in the amount of time required to process information (decreased information processing speed). For each experiment within a scenario, forager populations contained 0%, 25%, 50%, 75%, or 100% of cognitively impaired individuals, which reflects the range of frequencies of impaired individuals in previous studies on sublethal effects of pesticide and pathogen exposure on bumblebees. Each type of cognitive impairment was tested separately for each scenario. The starting populations were 100 bumblebees and either 4000 (scenario 1) or 1000 (scenarios 2 and 3) individuals of each plant species. For each experiment, we ran 100 replicates over 20 virtual seasons. Specific model variable settings used in each scenario are provided in Chapter 2.

## Effects on bee abundance only (Scenario 1)

We tested for sublethal stressor effects on bee populations when plant abundance and diversity were held constant (Table 9, bee abundance scenario). This scenario simulated an ecological condition in which bee foraging decisions had a minimal influence on the persistence of plant populations. Each plant species was assigned a different mean reward level so that foragers specializing on the species with the greatest mean reward maximized reward intake by the colony.

### Effects on plant diversity only (Scenario 2)

We tested for sublethal stressor effects on plant diversity when bee (pollinator) abundance was held constant (Table 9, plant diversity scenario). This scenario simulated an ecological condition in which stressor exposure altered bee foraging behavior, but not to an extent that reduced bee abundance. Plant population size in this scenario depended directly on the foraging decisions of bees (i.e., degree of conspecific and heterospecific pollen transfer). To allow increases in plant numbers over time and persistence of all plant species, we reduced initial plant population sizes and assigned the same mean (min, max) floral reward level to all species.

### Effects on system stability (Scenario 3)

We tested for sublethal stressor effects on the integrity of the entire system by allowing bee and plant populations to change over time in response to one another (Table 9, system stability scenario). This scenario simulated an ecological condition in which stressor effects on one side of the interaction (the pollinator, in this case) had the potential to influence population dynamics on both sides via their mutual feedback. Starting bee and plant population sizes and floral reward properties were the same as for scenario 2.

## Results

### Simulated effects of sublethal stressors on bee abundance (Scenario 1)

Under the control condition (0% impaired bees), bee population size increased slightly over 20 virtual seasons when plant populations were held constant (Figure 8). Bee population decline was observed when even 25% of the forager population had decreased memory capacity (Figure 8a) or decreased processing speed (Figure 8b), and the rate and magnitude of decline increased with an increase in the percentage of impaired individuals in the foraging population. For example, population decreases for bees with impaired memory capacity ranged from 26% (25% impaired) to 73% (100% impaired) after 10 virtual seasons. Slightly greater declines were observed when a percentage of bees had decreased processing speed; population decreases ranged from 33% (25% impairment) to 83% (100% impairment) over the same period.

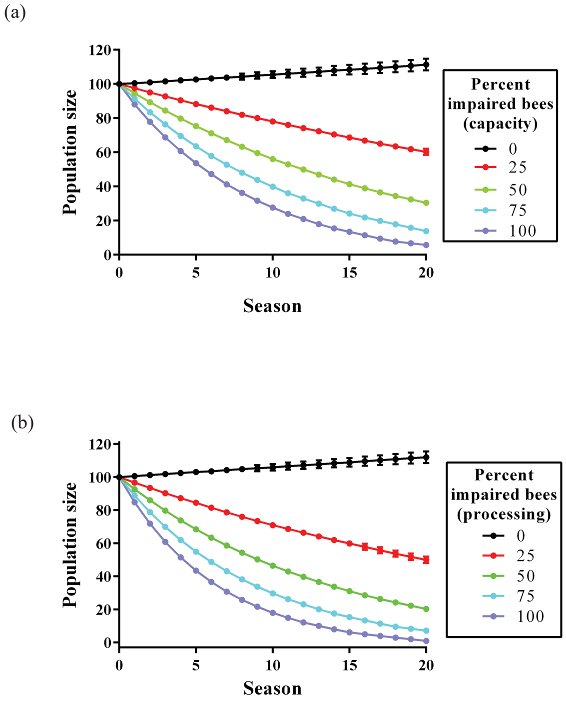

*Figure 8: Model simulations of sublethal stressor effects on bee population size with constant plant abundance and diversity (scenario 1). Mean (SD) number of bees in the system over 20 virtual seasons under decreasing (a) memory capacity and (b) processing speed (impaired, cognitive impairment; error bars, SD). One hundred simulations were run for each experimental condition.*

### Simulated effects of sublethal stressors on plant diversity (Scenario 2)

Under control conditions, a constant forager population of 100 unimpaired individuals maintained populations of all four plant species for 20 generations (Figure 9). Plant diversity decreased as the percentage of the bee pollinator population with decreased memory capacity increased (Figure 9a) due to increases in interspecific pollen transfer. Pollinator population impairment frequencies of 75% and 100% had the greatest effect on plant diversity over 20 virtual seasons and drove the complete loss of three out of four plant species for most simulation runs. The particular plant species lost was random because mean floral reward was the same for all species. Species loss resulted from small differences in plant species populations that arose stochastically and were amplified when impaired bees visited species at random over multiple virtual seasons. Thus, temporally, the dynamics of plant diversity loss followed a sigmoidal rather than linear function for all levels of bee impairment; all four plant species were maintained for the first six virtual seasons followed by a period of rapid species loss in the subsequent seven virtual seasons. In contrast, reductions in processing speed had a minimal effect on plant diversity (Figure 9b). The probability of losing a plant species in a mixed plant community therefore differed depending on the type and frequency of behavioral change in the pollinator population.

### Simulated effects of sublethal stressors on system stability (Scenario 3)

Under the control condition over 20 virtual seasons (0% of bees impaired), bee abundance increased, plant populations reached system carrying capacity, and plant diversity was maintained (Figure 10). For sublethal stressors affecting memory capacity (Figure 10a–c), bee and plant populations increased for all but the experiments with 100% impaired bees. However, as the proportion of impaired bees increased, so did heterospecific pollen transfer, resulting in reduced seed set and the eventual loss of plant species (Figure 10c). Under the 50% impaired condition, heterospecific pollen transfer increases alone were sufficient to drive the loss of at least one plant species in the 20 virtual seasons; most simulation runs showed a loss of two species. Loss of one or more species gave other species a better chance to be correctly pollinated by chance by memory-impaired bees; thus, overall plant population levels often reached the carrying capacity of the system, but with much less species diversity. Under 75% and 100% bee impairment conditions, all simulation runs ended with either a single plant species or complete system collapse.

For sublethal stressors decreasing information processing speed (Figure 10d–f), bee population size increased at a slower rate relative to the control at impairment frequencies of 25%, and effects on plant abundance and diversity were small. However, bee abundance decreased in the 50% impaired condition (Figure 10d), causing a corresponding decrease in plant populations and a loss of plant diversity (Figure 10e and Figure 10f) over time. Thus, unlike experimental runs involving decreased memory capacity, plant species declined due to the lack of bee pollinators in the system (pollen limitation) rather than reduced pollen transfer efficiency. Such negative feedback effects were much greater in the 75% and 100% impairment conditions and led to the eventual crash of the entire system in some simulation runs.

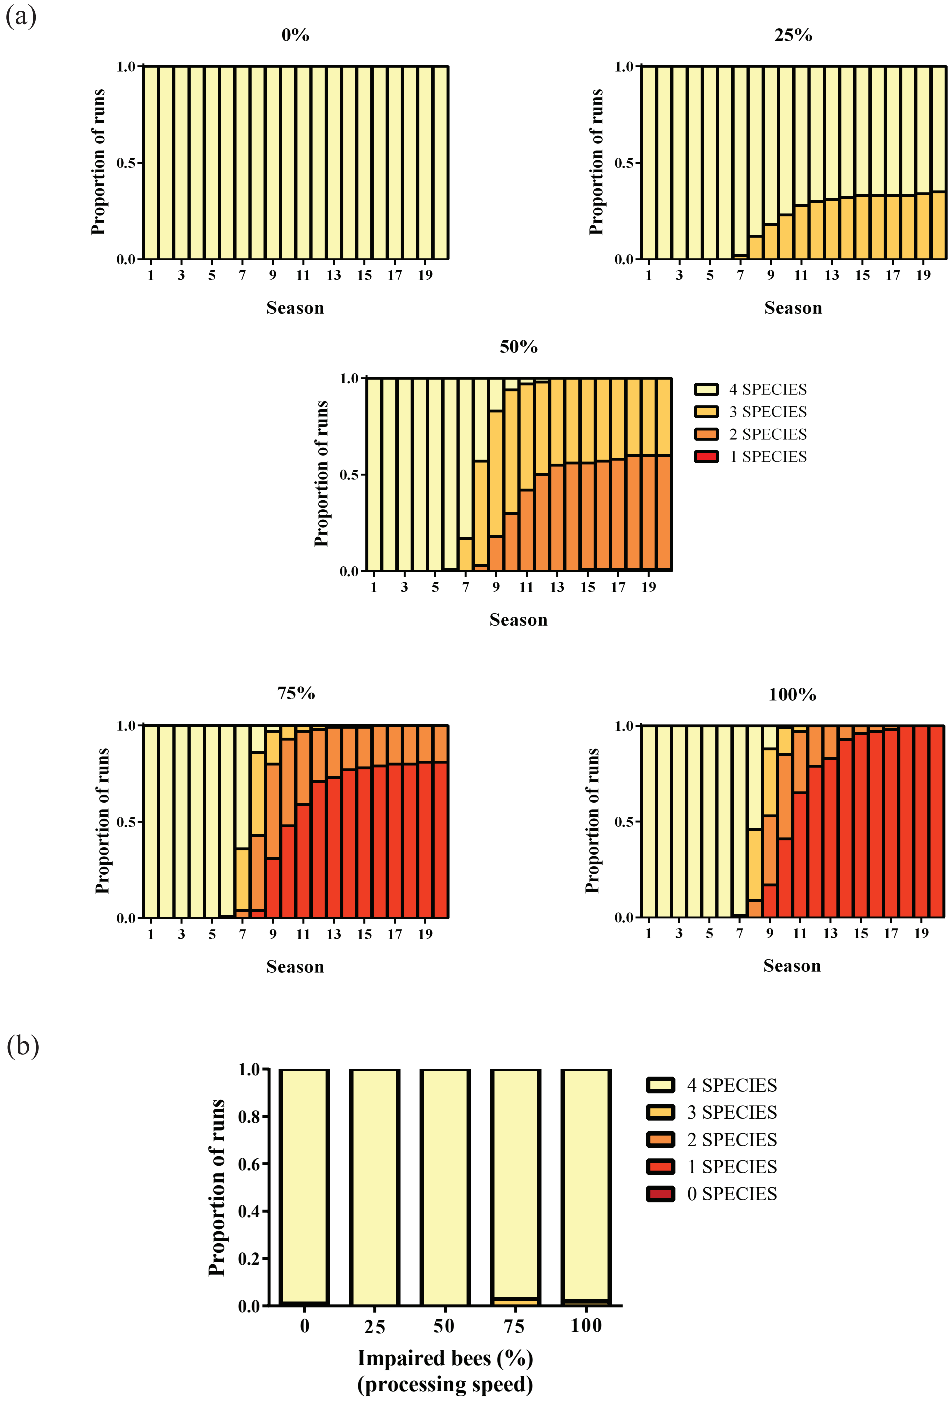

*Figure 9: Model simulations of sublethal stressor effects on plant diversity with constant bee population size (scenario 2). Proportion of simulation runs with 4, 3, 2, 1, or 0 plants species remaining at the end of each season: 0–100% of the bee population with (a) decreased memory capacity and (b) decreased processing speed. One hundred simulations were run for each experimental condition.*

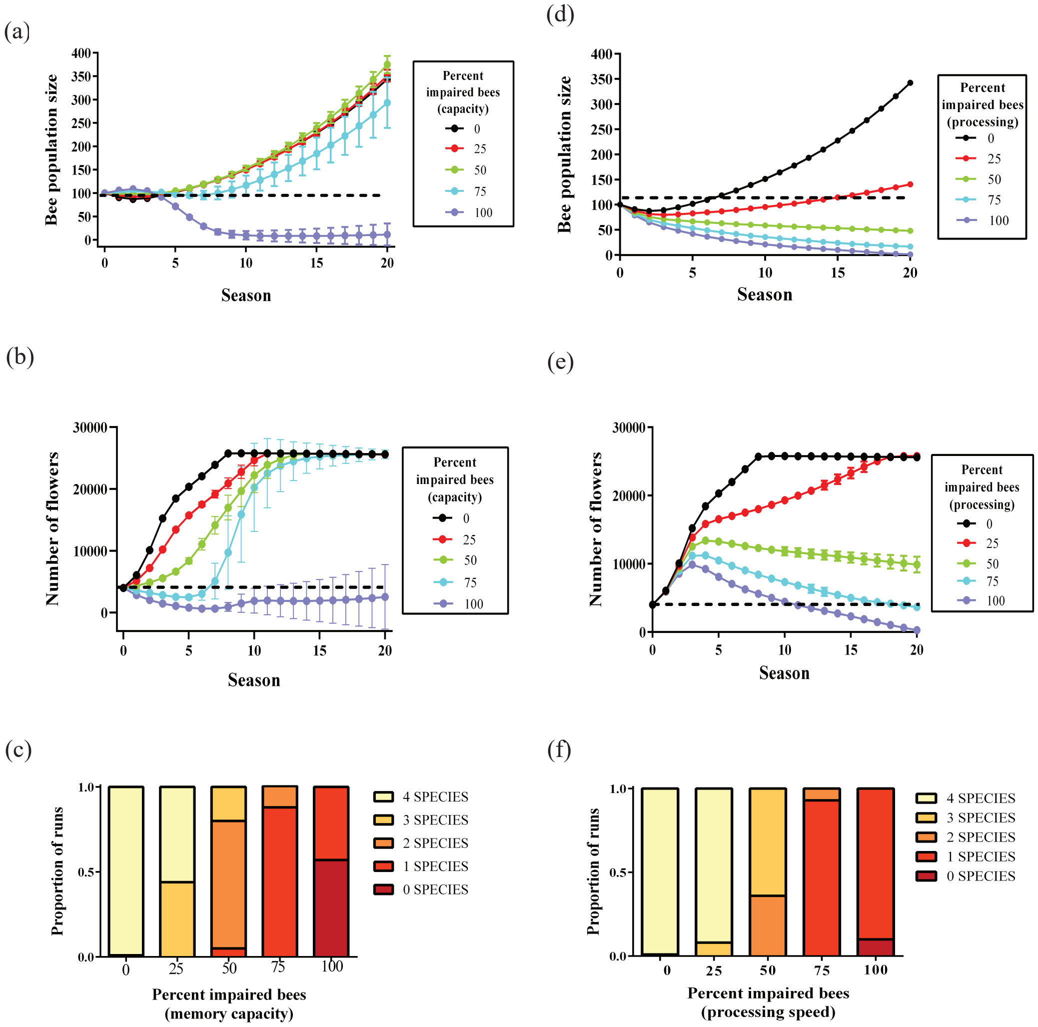

*Figure 10: Model simulations of sublethal stressor effects on bee–plant population dynamics (scenario 3). Mean (SD) of (a) number of bees and (b) number of rewarding flowers in the system over 20 seasons under decreasing memory capacity; (c) proportion of simulation runs with 4, 3, 2, 1, or 0 plants species after 20 seasons under decreasing memory capacity; (d) mean (SD) number of bees and (e) rewarding flowers in the system over 20 seasons under decreasing processing speed; and (f) proportion of simulation runs with 4, 3, 2, 1, or 0 plant species remaining at the end of 20 seasons under decreasing processing speed. (One hundred simulation runs for each experimental condition; carrying capacity, 25,921 plants [which provided an upper bound on bee abundance]). Test populations contained 0%, 25%, 50%, 75%, or 100% of foragers with cognitive impairment.*

## Discussion

Pollinator foraging decisions play an important role in maintaining the integrity of pollination systems. Past studies have shown that acute and sublethal exposure to a wide range of pesticides can impair the ability of pollinators to make adaptive decisions[37]. However, the vast majority of these studies have been on restrained individuals in a laboratory setting (e.g., the proboscis extension reflex paradigm), making the conservation implications of sublethal pesticide effects on plant-pollinator dynamics in the wild difficult to ascertain. Similarly, there is substantial evidence that pathogenic infection can impair pollinator decision-making[35] and in doing so reduce colony resource acquisition by increasing the amount of time required for foragers to process floral information[36,39]. Using an ABM approach, we found that sublethal effects on pollinator decision-making processes at the magnitude and exposure frequency reported in these previous studies have the potential to scale up and drive the decline and loss of bee species from plant–pollinator systems. Because our model is of a simplified pollination system, we cannot use it to make numerically precise predictions; however, our results are consistent with time frames for declines reported for natural pollinator populations. In addition, our study is the first, to our knowledge, to demonstrate that anthropogenic impacts on pollinator decision-making have the potential to drive the loss of species from plant communities without a parallel decrease in pollinator abundance.

### Bee population effects

Our simulations indicated that the impact of decreased memory capacity and information processing speed of bumblebee foragers on rates of population decline varied with the proportion of affected individuals in the population and with floral resource conditions. For example, forager memory impairment caused populations to decline rapidly when average reward level differed among available plant species (scenario 1), but population numbers actually increased when reward level was the same among species (scenario 3) at forager impairment prevalence of 75% or less. This finding was not unexpected, given that memory provided foragers with a clear adaptive advantage in scenario 1 (colony reward intake was greatest when foragers specialized on the most rewarding flowers) but not in scenario 3 (random foraging yielded the greatest energetic return). Decreased memory capacity did, however, drive bee decline in scenario 3 under the 100% forager impairment condition due to the increase in heterospecific pollen transfer, which drove the decline and eventual loss of plant species (floral resource availability). In contrast, decreased information processing speed (i.e., longer decision and flower-handling times) drove bee decline under both floral resource conditions. Thus, sublethal stressors could drive bee population decline in foraging habitats where memory demands are low or even absent, such as agricultural and urban landscapes dominated by a single plant species, if anthropogenic stressors with effects on information processing speed are present. Given that pesticides[146] and pathogens[144,145] known to impair bee cognition are both present in such landscapes, this is a reasonable possibility.

Our findings, in combination with the results of previous studies, indicated that sublethal stressor effects on bee cognition and behavior pose a significant threat to the abundance and diversity of wild bees. Past studies have shown that infection rates in wild bumblebees can be as high as 82% for naturally occurring pathogens such as *Crithidia bombi*[145] and as high as 40% for non-native pathogens such as *Nosema ceranae*[148]. Studies of wild bee populations in agricultural areas also show sublethal concentrations of insecticides such as clothianidin, thiamethoxam, and bifenthrin in 24–46% of individuals[149]. Our model represents a closed ecosystem with a simplified virtual season; thus, exact predictions of expected rate of decline for a particular species in the real world were not possible. However, the model predicted that cognitively impaired foragers at these frequencies would cause a 50% reduction in bee abundance in 4–6 years, which is consistent with declines reported for the endangered rusty patched bumblebee

(*Bombus affinis*) and other bumblebee species at risk in North America[14]. It is important to note that we used cognitive impairment levels based on data for a limited number of common, managed bumblebee species (*B. impatiens* and *B. terrestris*), which are likely to be less sensitive to sublethal stressors than wild species at risk. More comparative data on how the magnitude of stressor-induced decreases in cognitive performance varies among bumblebee species are needed to test this possibility. The incorporation of multispecies data on stressor-induced behavioral changes into our model could also be used to investigate how anthropogenic disturbances alter competition in pollinator communities, ultimately causing loss of some species and dominance of others[150]. Further development of our model to investigate how other forms of anthropogenic disturbance, such as habitat modification and fragmentation, affect the foraging success of bees and other pollinators could also accelerate conservation efforts to identify, protect, and restore threatened species.

### Plant community effects

In addition to driving pollinator decline, our results showed that changes to pollinator behavior caused by stressor-induced cognitive impairment could have profound negative effects on the structure and diversity of flowering plant communities by affecting either the quality or quantity of pollen transferred among individual plants. In scenarios 2 and 3 under control conditions, populations of all four plant species were stable over 20 virtual seasons because individual bees exhibited floral constancy (i.e., individual foragers temporarily specialize on flowers of one plant species despite the availability of equally rewarding flowers of other plant species[67]. Flower constancy increases plant reproductive success by reducing heterospecific pollen transfer.

Our results indicated that even a marginal stressor-induced decrease in flower constancy due to impairment of memory capacity was sufficient to cause the decline and loss of species from plant communities in scenarios 2 and 3 (Figure 9a and Figure 10a,c). In contrast, a decrease in information processing speed had no effect on bee constancy levels and therefore had no effect on plant abundance and diversity in scenario 2 (Figure 9b). Decreased processing speed resulted in plant species loss in scenario 3 (Figure 10d–f); however, it was driven by the decline of bee populations (pollen quantity) rather than flower constancy (pollen quality). Previous modeling approaches to plant species extinctions in pollination networks were based on the assumption that pollinator loss is necessary to initiate plant extinction events[151,152]. However, we found that changes to pollinator behavior alone could cause the loss of plant species even when pollinator populations remained stable (Figure 9a) or even increased over time (Figure 10a–c). In fact, we found behaviorally driven single and multiple species extinction events in plants occurred in as few as seven virtual seasons and coextinction events (bees and plants) occurred in as few as nine virtual seasons.

Based on these results, we expect the greatest risk of plant extinction due to changes in pollinator behavior to occur in mixed plant communities containing species with high levels of pollinator specificity, reproductive dependence on animal pollination, and demographic dependence on seed production[56]. From the perspective of biodiversity conservation, changes to pollinator behavior resulting in reduced pollination efficiency can negatively affect secondary wildlife consumers through reductions in the availability of seeds and fruit even when plant species have a low dependence on pollinators or can compensate for decreased seed production[153,154]. Interestingly, when plant species diversity in our model dropped to just one or two species due to memory capacity impairment, populations of bees and those plant species began to increase because of the reduced chance of heterospecific pollen transfer. Thus, pollinator behavior can provide a better proxy for plant extinction risk than pollinator abundance in some cases and therefore must be considered when developing effective conservation and management strategies for threatened pollination systems.

### ABM and the conservation of plant–pollinator systems

Our findings highlight the importance of using an ABM approach to understand how anthropogenic disturbance influences the structure and dynamics of pollination networks through changes to pollinator behavior. Mathematical network models often use the degree of interactions among species to predict extinction scenarios[12]. However, it is well documented that the behavior of individuals within a species is often far more specialized than the average behavior over the population[155]. For example, flower constancy in bumblebees is observed at the level of individuals, even in generalist species[59,156]. In our virtual system, flower constancy was an emergent property resulting from interactions between memory limitations in bees and floral resource variability among plants. Failing to account for the behavior of individuals can result in underestimates of the likelihood of extinction of a species[157,158]. More recent stochastic extinction models recognize that interaction strengths can vary and that plant–pollinator mutualisms are more important in predicting species loss than interaction strength alone[151,152]. Although network models have become increasingly complex and small scale, they still lack the fundamental individual mechanisms that determine the tipping points between population persistence and decline[101]. However, network models include data that are critically relevant to extinction prediction. In fact, inclusion of field data on variation in floral preference among bumblebee species and mating strategies among plant species in natural systems, as well as seasonal variations in pollinator and plant populations, will be an important future step toward improving the ability of our model to predict the long-term effects of anthropogenic disturbance on biodiversity and ecosystem health.

# Chapter 4: Response to Change & Memory Dynamics in Bumblebees

## Introduction

How do bumblebees convert recent experiences into decisions? Do bees simply use the most recent decision to make a choice? Do they 'average' recent outcomes to approximate an expected value, or do they remember probabilities of reward? Do they prioritize recent outcomes more heavily as uncertainty increases? How do bees manage information and expectations in dynamic and variable environments? Understanding how bumblebees utilize experiences is key to predicting their ability to adapt in rapidly changing environments.

It is well-established that bumblebees must constantly learn about their environment, form associations, and make discriminations as they forage for resources during their limited life-cyle[52,58,159,160]. Food sources in nature vary both temporally and spatially, making the ability to keep track of a variable and/or changing environment critical for decisions as to when and where to forage[72,160]. In a changing environment where the most rewarding option changes over time, the ability to recognize change in the environment and update associations based on new information is adaptive[159]. In a variable environment where the probability of receiving a reward from an option is uncertain, bees must use memory to determine which option is actually the most profitable choice[58]. How bees achieve this with a small and relatively simple brain is extensively debated[83,161,162]. Memory plays an important role, as evidenced by decreased foraging abilities when memory is impaired[38,39]. However, the link between memory and decision-making in foraging bumblebees is not fully understood. Making decisions based on past experiences is detrimental in a rapidly changing environment where old memories are no longer reflective of the present environment but making decisions based on only the most recent experiences in a highly variable environment will result in suboptimal decisions due to an inaccurate estimate of the environment. While limiting the use of memory could be advantageous in changing environments, the same limitations could hinder foragers in variable environments[163]. Since bees can encounter both variable and changing environments in the wild, they must develop ways of adapting to both types of conditions. Characterizing how bumblebees can adapt to rapidly-changing and variable environments can provide insight into the reasons why some bumblebees adapt quickly to human-induced environmental change[164] (HIREC) while others fail to adapt[11,18,164].

## Behavioral Flexibility

While memory dynamics and decision-making are different from behavioral flexibility, it is important to understand previous research on behavioral flexibility when considering the cognitive abilities of bumblebees. Bees are often the subject of research in animal cognition due to their aptitude for complex cognitive tasks[92,165]. The term "behavioral flexibility" is often used interchangeably with "cognitive flexibility" to describe the ability of individuals to adapt to certain environmental conditions[52,166,167]. However, such a definition is too general to be helpful in describing the cognitive abilities of bumblebees and resulted in considerable debate around the meaning of the term and how it can be tested[168,169]. Instead, I focus on a subset of the definition rooted in psychology: the ability of individuals to make decisions influenced by various schedules of rewards[168,170]. Reward schedules where the probability of receiving a reward is uncertain but consistent simulate variable environments while reward schedules where options are swapped simulate changing environments. How individual bumblebees adapt to changes in rewards may be informative for predicting both their successes and exposures to risk caused by human-induced rapid environmental change[171–173]. While the behavioral flexibility of bumblebees and their decision-making abilities are well documented and have even inspired novel algorithmic approaches to problems in engineering[174,175], the link between observed behavioral flexibility and the underlying strategies and mechanisms of bumblebee decision-making and memory remains elusive.

In a traditional laboratory context, a bumblebee is considered behaviorally flexible if it can successfully learn new information while influenced by past information[159]. A popular framework for measuring such flexibility in individuals is the *reversal learning* experimental paradigm[170,176,177]. In reversal learning, an individual is trained on a particular discrimination test – a task where an individual learns to distinguish between two options - and then, once a criterion level of accuracy is reached, the contingencies are changed or reversed[178]. For example, a bumblebee may be presented with a blue and an orange test tube and is trained such that the blue tube contains a sucrose reward while the orange tube contains no reward. This scenario is repeated until the bee reaches a predetermined level of accuracy in choosing the rewarding-colored tube. After the bee reaches the performance criterion, the relationship is reversed such that the orange tube is rewarding while the blue tube is not. A bee that displays high behavioral flexibility is one that takes few trials to reach the performance criterion by reversing the previously reinforced behavior. This experimental paradigm tests associative learning, a process often considered simple but requiring relatively complex neural circuitry[162,179]. Such measures of behavioral flexibility have historically been useful in examining the mechanisms of learning and memory[74,97,180] and thus represent a useful avenue for investigating those mechanisms through simulation modeling.

Although reversal learning is useful for exploring behavioral responses, such experiments typically do not include variability in rewards, or the possibility of an option being rewarding only with a certain probability (partial reinforcement). Here, we draw inspiration from human psychology using experiments on decisions from experience. In these experiments, the human subject is given a binary choice task and receives feedback about their choice[181–183]. In some versions, one option changes the probability of reward over time while the other option provides a 'safe' but low guaranteed reward (safe vs risky). In other versions, one option has a higher expected outcome than the other, but neither option provides a reward with 100% probability (risky vs risky). Halfway through the experiment, the options switch, requiring participants to adapt their choices in order to maximize their outcome[184]. Both versions of decisions from experience experiments test how individuals adapt to change, but recent work demonstrates that risk preferences and the direction of change influence adaption to change in the safe vs risky version of these experiments. Specifically, individuals were more likely to switch away from a highly rewarding risky option that became less rewarding than they were to switch to a risky option that became more rewarding[181]. Comparisons with a similar risky vs risky experiment showed that the effect of the direction of changed disappeared when the riskiness of both options was the same, demonstrating that risk aversion – the preference for more consistent but less rewarding options – can influence an individual's response to change[184]. Experiments with risky options can isolate potential mechanisms for change from the effects of risk aversion.

These experiments with human subjects have identified several potential mechanisms that are involved in the detection of and adaptation to change[181]. The over-reliance on initial experiences – often referred to as *stickiness* – occurs when a subject fails to recognize that rewards have changed and continues to make the same decisions. While this can be beneficial in certain highly variable environments that do not change overall, this stickiness effect is often seen as detrimental for adaption to change. In contrast, the over-reliance on the most recent experiences – *recency* – is a form of cognitive bias that favors recent experiences over old ones. Individuals influenced by recency bias make decisions that ignore long-term patterns, something that is disadvantageous in predictable environments but beneficial in rapidly changing environments where old patterns may no longer be relevant[181,182]. Overall, human subjects usually demonstrate difficulty adapting to change and tend to over-value initial experiences.

Additionally, altering probabilistic reward schedules can also answer questions about the effects of partial reinforcement in bumblebees. If reversals require unlearning of initial responses, partial reinforcement should result in bees learning the reversal faster. However, similar experiments in rats found that partial reinforcement led to a weaker learned response, with rats taking much longer to reverse their initial preferences when the reward schedule frequency was not 100%/0%[185,186]. Learning how bumblebees respond to partial reinforcement is useful for understanding how past and present experiences influence decision-making.

## Past Efforts to Elucidate Bumblebee Strategies and Memory

A survey of the literature regarding bumblebee decision-making reveals a variety of theoretical models. These models range anywhere from simple rule-based models such as “Win-Stay Lose-Switch” (WSLS) to Optimal Foraging. I cover these models in detail in a later section, but generally most existing models either use statistical calculations of profitability (optimal foraging, optimal diet, and Bayesian derived models)[139,187–189], or use simple heuristics that replicate observed behavior[190,191]. These models are validated through comparison with laboratory or field observations of bumblebee behavior and demonstrate that bumblebees do have the behavioral flexibility to adapt to changes in their environment. However, they do not help us understand how bumblebees adapt to more realistic environments, where rewards may be variable as well as changing over time. To better understand how bumblebee memory functions mechanistically, we wanted to determine whether some combination of memory capacity and decision-making strategy might be optimal for these more complex and realistic environments. None of the models mentioned in Table 1 allow us to simulate foraging behavior in variable and changing environments, so we consider how we can consider the role of memory and decision-making in these environments using the previously established SimBee model.

Modeling memory in an accurate way is important for mechanistic models. While optimal foraging models can reproduce foraging behavior under certain conditions, researchers recognize that bumblebees are not optimal foragers and make choices that are inconsistent with optimal foraging models that assume global knowledge of the environment[31,58,138]. Indeed, more recent research recognizes that the limited working memory of bumblebees result in suboptimal choices[188]. Based on studies of risk-sensitive foraging, it is known that bumblebees can recognize changes in their environment and switch their foraging strategy[58]. However, bumblebees are clearly not optimal foragers and exhibit a delay before switching away from non-rewarding flowers. This delay is a key indicator of how much foraging bees weigh their old memories of their environment against new memories.

Despite the evidence that bumblebee memory and decision-making is influenced by limited memory, existing agent-based mechanistic models do not consider the memory limitations. In the past 10 years, only two relevant ABMs have been published: Qu et. al 2013 and Becher et. al 2018[100,109]. The Qu 2013 model implements reinforcement learning for its agents to inform patch-based decisions but does not limit the memory of bees. The Becher 2018 ABM does not consider individual bumblebee memory and treats foragers as cohorts of bees that make decisions at the patch level. These models are summarized in Table 10. Other memory models can be found that are not covered here, but they lack memory components and methods of detecting change in the environment. The lack of a mechanistic model for individual bumblebee memory and decision-making is a serious knowledge-gap that needs to be filled.

To address the knowledge gap in bumblebee memory and decision-making, we extend SimBee to simulate how bumblebees with different models of memory and decision-making forage in changing and variable environments. We expand on existing work with reversal learning experiments with bumblebees and decision from experience experiments with humans by designing three experimental scenarios to investigate behavioral responses to three different types of change. Based on the decision from experience experiments described earlier, we use similar paradigms to determine what decision-making strategies bumblebees use, as well as what memory capacity enables adaption to change. To evaluate these questions, we compare the simulated behavior of different models of memory and decision-making to preliminary data, as well as determining which models result in optimal performance.

*Table 10: A collection of memory models found in the literature that can be applied to bumblebee behavior.*

| **Paper** | **Memory Model** | **Decision-Making Strategy** |
|---|---|---|
| Bélisle and Cresswell, 1997[192] | Sliding Memory model | Markov Chain Probabilities |
| Zhang and Hui, 2014[188] | Optimal diet model with memory component | Maximize reward/time based with knowledge limited to past experiences. |
| Zhang et. al 2005[95] | Exponential Decay | Probability of recalling the correct decision decreases exponentially |
| Qu et. al 2013[109] | Reinforcement Learning | Maximize expected value |
| Becher et. al 2018[100] | Optimal foraging | Maximize fitness |
| Keasar 2002[193] | Rescorla-Wagner Linear Learning Rule | Maximize expected value |

## Bumblebee Strategies

To determine what memory dynamics and capacity might best explain bumblebee decision-making behavior in scenarios with both change and variability, we considered a number of models derived from past literature and describe common approaches below.

## Optimal Foraging

Researchers have been working on models of bumblebee memory and decision-making since the 1970s. Most of these early approaches focused on mathematically determining the optimal behavior of a bumblebee forager through predator-prey models. The predator-prey model assumes an environment with randomly distributed prey of various types. Each prey type has an associated handling time, reward, and encounter rate. When the predator encounters prey, it must decide between spending time to handle the prey or continue searching. If the predator knows these values for every prey type in the environment, it can make the optimal decision to either accept or reject the prey. When applied to pollinators, the forager becomes the predator seeking to extract nectar rewards from stationary plants. The idea that natural selection would favor foraging strategies that maximize reward per time (or minimize effort per reward) gave birth to optimal foraging theory, an early mathematical model of foraging behavior[139,140]. This model is based on the key assumption that animals forage in a way that maximizes their fitness, hence the name *optimal foraging theory*[194]. The fitness of an individual is measured by some currency that is usually in the form of energy / time. While there are numerous forms of the optimal foraging model depending on the exact scenario, the maximization function almost always appears in the form:

$$Mean\ net\ rate\ of\ energy\ gain = \frac{Mean\ energy\ gained - Mean\ energy\ spent\ gathering}{Time\ spent\ traveling + Time\ spent\ gathering}$$

*Equation 6*

Given perfect knowledge of the environment, this function can be used to determine the optimal decision for a foraging bumblebee. Some studies found that in the wild, bumblebees can match the predicted behavior of the optimal foraging model in terms of distance traveled and direction of departure chosen[140,194]; however, validating such models in the field based on energy maximization was not feasible. Eventually, more detailed studies determined that the disconnect between the complexity of natural environments and the assumptions made in the optimal foraging model indicated that more complex models of decision-making were needed[140]. While the optimal foraging model is still used as an elegant model for optimized behavior given energy and time constraints, real bumblebees almost never match these optimal behaviors. While we still assume foragers aim to maximize their energy intake, for these reasons and to avoid the complexity of adding travel and gathering time to SimBee, we do not consider optimal foraging related models in this research.

### Heuristics

Heuristics are simple models that allow individuals to quickly determine a feasible solution while ignoring some of the available information[195–197]. A key characteristic of heuristic approaches is that they seek an acceptable solution and not the optimal one which is associated with more complex models of decision-making[198]. Compared to more complex models, heuristic models are advantageous in terms of reduced time, information, and energy required for a decision while achieving similar accuracy to their counterparts. On the other hand, heuristics inherently introduce systematic errors, biases, and deviations from the optimal solution[198,199]. Compared to memoryless models, heuristics represent a vast improvement over random decision-making with a relatively low added cost of information. This provides a class of models that fit the suboptimal but sufficient results often seen in human and animal behavioral studies[200]. A key component of heuristic models is the limited memory. Only considering a small number of past experiences is the central component of heuristics that make them suboptimal models but more accurate representations of decision-making. As discussed in chapter 1, bumblebee memory – like human memory – is known to be limited, and we believe decision-making models should reflect this. We include several heuristics chosen from a study on human decisions from experience as potential models for bumblebee decision-making[201]. The heuristics are described in Table 11. It is important to note that while these heuristics have been used in controlled studies, some, such as the lexicographic model, may not translate well to real-world experiences where rewards are variable and continuous without modification.

## Weighted Memory Models

The idea that old experiences are given less weight than new experiences when making decisions has spurred debate in many areas related to cognition[176]. In human psychology, this idea is often referred to as the "forgetting curve", a non-linear function that relates the probability of memory retention to the delay between acquisition and use[202]. While the forgetting curve has been a central question to human psychology for a very long time, there is lack of consensus over the correct form and parameterization of the function[203]. Despite this uncertainty, there does seem to be general agreement that the general expression for the forgetting curve should be:

$$R(t) = a + (1 - a) \times b \times P(t)$$

Where R(t) is the probability of retention, P(t) is a nonlinear function where P(0) = 1 and P(t) approaches zero for large values of t, and the parameters *a* and *b* are bounded between zero and one. As the nature of the forgetting curve can vary drastically depending on the form of P(t), the form of the forgetting curve is often chosen based on the scenario being modeled.

Drawing inspiration from the previous collaboration on vehicular communication (see Appendix S2: Bumblebee-Inspired Vehicular Communication Algorithms), we simplify this general equation by letting $a$=0 and b=1 and using either a logarithmic function or exponential function for P(t). Therefore we can derive an exponential forgetting curve as:

$$R(t) = e^{-\propto t}$$

*Equation 7*

And a logarithmic forgetting curve as:

$$R(t) = \propto \ln t$$

*Equation 8*

If we consider R(t) as the weight of memory in calculating a weighted mean, we can use these equations to generate decisions based on past experiences. As long as $0 < \alpha < 1$, then older memories receive less weight than newer memories. Thus, we transform the general forgetting curve into a weighted memory model. Both models weight the oldest memory the least and the newest memory the most but differ in how they weight memories in between as shown in Figure 11.

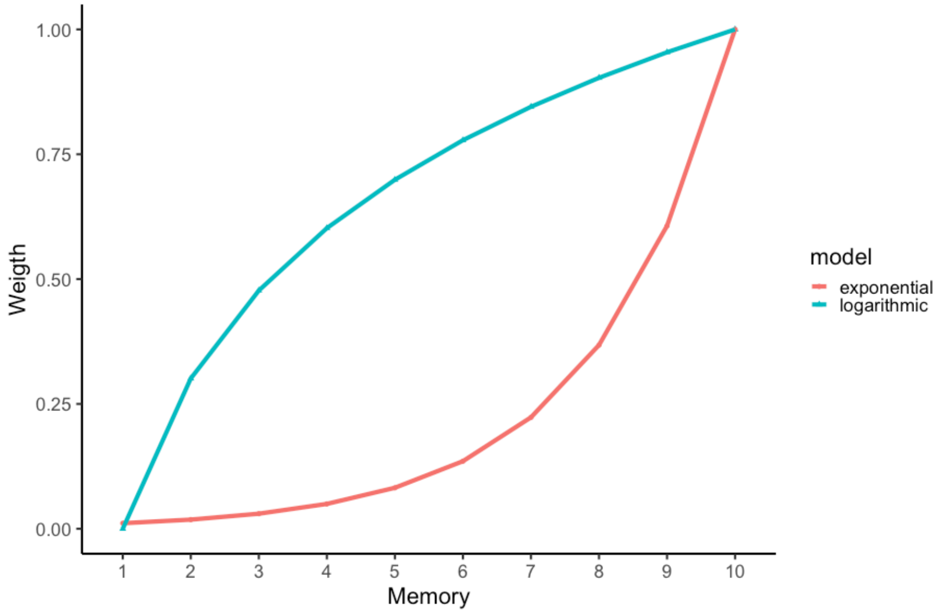

*Figure 11: The weights of the exponential and logarithmic memory models plotted as a function of memory recency.*

The exponential and logarithmic weighted memory models provide a more complex decision-making algorithm than the common heuristics discussed previously. This level of complexity may be required to optimally exploit certain complex environments and provide an additional comparison for modeling bumblebee decision-making.

## Reinforcement Learning

Reinforcement learning (RL) is learning how to maximize a numerical reward for a given situation. Through trial and error, the learner must discover which actions yield the most reward. The key aspect of RL is the trade-off between exploration and exploitation. To maximize reward gain, a reinforcement learning agent must prefer actions that it has found to be effective in the past. However, to discover those actions the agent must also try actions it has not selected recently. The dilemma is that neither exploration nor exploitation can be used exclusively without failing to maximize the reward gained. The fact that RL incorporates this dilemma makes this type of problem-solving method popular in a variety of fields. In Ecology, a special case of RL is often referred to as the Rescorla-Wagner learning rule[204] and has been used to model bumblebee decisions in a two-armed bandit experiment[193]. In a multi-armed bandit situation, a bee must choose between multiple flower options to visit when the profitability of those options is not known in advance. Similar to bees foraging in a field, the bees in this type of experiment face the same exploration-exploitation dilemma as they attempt to maximize their reward intake. The general form of the Rescorla-Wagner learning rule is referred to in psychology as the delta learning rule[181,205,206], in reference to updating expectations based on the difference between the previous expectation and current reward. The delta learning rule has been used extensively in modeling experience-based decision-making tasks[181,207] and takes the general form:

$$E_{j,t} = E_{j,t-1} + \delta_{j,t}\alpha(r_j - E_{j,t-1})$$

*Equation 9*

Here, the expected value $E$ of an option *j* on trial *t* is the sum of the previous trial $E$ adjusted by the prediction error $r_j - E_{j,t-1}$, which is the difference between the reward received on trial t and the expected reward after trial t-1. The adjustment is controlled by the updating parameter $\alpha$ that ranges between 0 and 1, where values close to 0 indicate weak recency effects and values close to 1 indicate strong recency effects. Consequently, the $\alpha$ parameter here serves the same function as the $\alpha$ parameter in the weighted memory models discussed previously. The $\delta$ parameter ensures that the delta rule only updates the expected value of the chosen option; the values of the unchosen options are unchanged until they are selected.

The decision in the RL delta learning rule model is implemented by a softmax choice rule with the single inverse-temperature parameter $\theta$ :

$$P(Choice = j) = \frac{e^{\theta E_j}}{\sum_{k=1}^{k} e^{\theta E_k}}$$

*Equation 10*

Values where $0 < \theta < 1$ will generate probability distributions that are more concentrated around the smallest input values (promoting exploration), while values of $\theta$ greater than 1 will create probability distributions concentrated around the larger input values (promoting exploitation). As $\theta \to \infty$, softmax converges to argmax.

## Methods

### Proposed Memory Models

Table 11 summarizes the different rules, heuristics, memory models, and strategies we implement in SimBee to determine which produces optimal bumblebee behavior in terms of reward gathered in specific environments. These strategies are designed for scenarios that are experimentally feasible so that future work can compare experimental results to the optimized behavior. These strategies can also be compared to experimental data to determine which strategy generates the most similar behavior.

*Table 11: Memory models implemented in SimBee. These decision-making strategies are possible explanations for how foraging bumblebees use past experiences to make decisions and adapt to changing environments.*

| MODEL | DESCRIPTION |
|---|---|
| **WIN-STAY LOSE-SWITCH[208] (WSLS)** | Step 1: If the previous decision returned a positive outcome, choose the same option. Otherwise choose a different option randomly. |

| | |
|---|---|
| **BASELINE** | Step 1: Calculate the arithmetic mean of all experienced outcomes within an option<br>Step 2: Choose the option with the higher mean |
| **BETTER THAN AVERAGE**[201] | Step 1: Calculate the grand average of all experienced outcomes from all options.<br>Step 2: For each option, count the number of outcomes equal to or above the grand average.<br>Step 3: Choose the option with the highest number of such outcomes |
| **LEXICOGRAPHIC**[201] | Step 1: Determine the most frequently experienced outcome of each option.<br>Step 2a: Choose the option with the highest most frequent outcome.<br>Step 2b: If both are equal, determine the second most frequent outcome of each option, and select the option with the highest (second most frequent) outcome. Proceed until a decision is reached |
| **TALLYING**[201] | Step 1: Give a tally mark to the option with (a) the higher minimum gain, (b) the higher maximum gain, (c) the lower experienced frequency of the minimum gain, and (d) the higher experienced frequency of the maximum gain.<br>Step 2: Choose the option with the higher number of tally marks |
| **LOGARITHMIC** | Step 1: Calculate the weighted mean of all outcomes within an option, where the weights follow a natural logarithmic function based on the time of occurrence of each outcome.<br>Step 2: Choose the option with the higher mean |
| **EXPONENTIAL** | Step 1: Calculate the weighted mean of all outcomes within an option, where the weights follow an exponential function<br>Step 2: Choose the option with the higher mean |
| **REINFORCEMENT LEARNING (RLDELTA)** | Step 1: After choosing an option, update the expect value of the option based on the difference between the expected outcome and actual outcome.<br>Step 2: Choose an option based on the weighted probability distribution of expected values using a softmax function (See Equation 10). |

## Modifications & Additions to SimBee

### Split Memory

In the published version of SimBee, memory is structured as a list of reward-plant species pairs, with 20 being the maximum number of pairs in the list (*max-memory-length*). This is implemented as a list of tuples called *visited-memory*, which we classify as a single-memory structure. However, the degree to which past memories interfere with new memories – especially those with different associations – is unknown[209,210]. To address this knowledge gap, we implement an alternative memory structure henceforth referred to as split-memory. Instead of a single list of tuples, split-memory consists of a separate list of reward values for each species of plant present in the simulation. An example of this with two species of plants is shown in Figure 12.

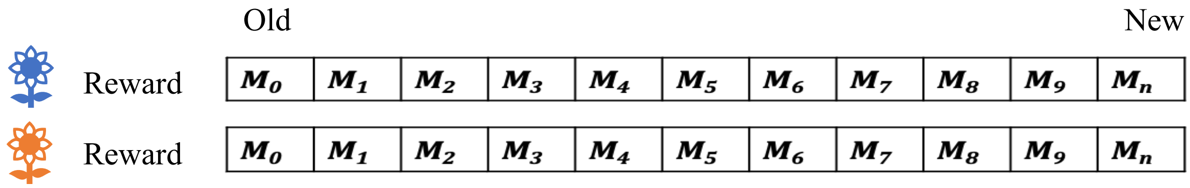

*Figure 12: Example of split-memory in SimBee. Experiences from each plant species are stored in a separate list rather than all together in a single list of tuples.*

Like the original implementation of memory in SimBee, once memory capacity is reached, each new reward added to the matching species list removes the oldest memory from that list. In all other regards, split-memory functions the same as single-memory.

### Sampling

Sampling mode in SimBee was a time-based behavior where *bee* agents sampled *plants* for 100 ticks at regular intervals. However, we do not consider sampling for the scenarios discussed in this chapter. The timeframe for these experiments is short enough and options limited such that sampling behavior as described in Chapter 3 should not occur. We remove the periodic sampling behavior and only ensure that virtual bees have their memories filled with reward outcomes of one for each option at the beginning of the foraging run, indicating to the bees that both options are equally rewarding.

## Choose-Plant

Numerous alterations were made to the first two steps of the *choose-plant* function to accommodate the implementation of new heuristics and memory-based decision-making modules.

### *Reinforcement Learning*

The *RLDelta* option under the interface chooser 'bee-memory-model' requires several deviations from other memory-based decision-making models. First, the memory size of the *bee* agent is set to one, as only the current expected values are saved for each *plant* species. Second, reinforcement learning is the only memory model that utilizes the *bee-softmax-theta* parameter, which must have a value. Finally, the *choose-plant* function follows a different set of steps to determine the correct choice, as explained below:

**Step 1: Calculate expected values for the reward of each plant species.**

Bee agents only store the expected value $E_j$ for each species of plant $j$ in memory. After collecting a reward $r_j$ for a plant, the bee updates the expected value in memory. The updating process after visit $v$ can be summarized as follows:

$$E_{j,v} = E_{j,v-1} + \alpha(r_j - E_{j,v-1})$$

Equation 11

Alpha ($\alpha$) is either set by the user or varied between 0 and 1 for parameter estimation.

**Step 2: Determine action probability for each species.**

The expected values $E_j$ for each species of plant in view are converted into action probabilities using the softmax function with temperature parameter $\theta$ as follows:

$$P(Choice = j) = \frac{e^{\theta E_j}}{\sum_{k=1}^{k} e^{\theta E_k}}$$

Equation 12

Where *k* is the number of plant species in view.

**Step 3: Choose species to visit.**

Given the probability distribution calculated in the previous step, randomly select a species of plant to visit. The probability of each option being picked is proportional to the weight given by the probability distribution. The process then continues as outlined in the SimBee ODD at step 4.

*Heuristics*

*Bee* agents using heuristics to make decisions utilize the *heuristic-decision* function to evaluate experiences in memory and return the correct choice. Each heuristic has its own helper function for ease of use.

## Plant Reward

In the published SimBee model, plant rewards are generated from a uniform-random distribution with a min and max value specified by the user. To address other research questions, three additional methods for plant reward generation are implemented in the SimBee model as described below.

### *Sequential*

The sequential refill method allows the user to specify a sequence of one or more reward values for each plant species. Sequential visits to a specific plant species will result in the rewards specified by the list provided, in that exact order. Once the list is exhausted, the sequence of rewards repeats indefinitely. This implementation is designed to mimic rewards in laboratory experiments.

### *Bounded Normal Distribution*

The bounded normal distribution method returns a reward generated by a standard normal distribution given a user specified mean and standard deviation. Values outside of the user-defined lower and upper bounds are discarded and another value is generated until a value that falls within the bounds is returned. Negative values can be generated if the lower-bound is less than zero, but negative reward values will not be accepted by *bee* agents.

### *Probability of Refill*

For the probability of refill method, the user specifies a reward value using the *plant-probability-refill* input and a value between zero and one inclusive for each plant species. These are independent probabilities of a visit to each plant species returning the reward specified in *plant-probability-refill*; otherwise, the plant returns a reward of zero.

### Reverse-environment

The function *reverse-environment* is designed for simulating reversal learning scenarios. When a user-specified number of visits are recorded, this function is called to switch the reward schedules of plants in the simulated environment. Using the same reward schedule setup as specified by the user in the interface, the reward schedules for plant species 1 and 2 are swapped, and the reward schedules for plant species 3 and 4 are swapped. The plants are then asked to regenerate their existing rewards. This guarantees that any subsequent visits will experience the reversed reward schedule.

### Fitting memory models / Parameter Estimation

Parameter estimation in NetLogo simulations was done using the BehaviorSearch tool designed by Stonedahl and Wilensky[211]. BehaviorSearch is an open-source cross-platform tool that provides several search algorithms and search-space representations/encodings and can be used to explore the parameter space of any ABM written in the NetLogo language. For comparisons with experimental data, the objective function was designed as the mean-squared error (MSE) between the experimental data and simulated data for the choices of a single bee agent in blocks of twenty visits. If experimental data was not used, the objective function maximized the amount of reward collected by an individual bee at the end of the simulated run. In both cases, the measure of fitness was collected at the final step. The genetic algorithm approach with default parameters was used with a MixedTypeChromosome encoding due the mixed data types of the model parameters[211]. The mean fitness of 100 replicates was used during each step of the process. Parameters that were optimized are summarized in Table 12.

*Table 12: Summary of parameters that are optimized using BehaviorSearch in each scenario. Not all parameters are used for every model; for instance, RLDelta does not optimize memory size but is the only model to use softmax theta. WSLS and Random Choice do not require parameter estimation as they do not utilize memory size, memory alpha, or softmax theta.*

| PARAMETER | DESCRIPTION | RANGE |
|---|---|---|
| **MEMORY MODEL** | Which memory model the bee uses. | See Table 11 for all memory models implemented in SimBee |
| **MEMORY SIZE** | Memory capacity of the bee, specifically the number of experiences for each species of plan the bee stores. | Heuristics and weighted memory models: [3,10]<br>WSLS and Random: NA<br>RLDelta: 1 |
| **MEMORY ALPHA** | Parameter in Weighted Memory models and RLDelta for recency bias. | Weighted memory models and RLDelta: [0,1] in increments of 0.05 |
| **SOFTMAX THETA** | Temperature parameter for RLDelta softmax function. | RLDelta: [0, 100] (continuous range) |

While the range of these parameters were limited primarily by their implementations, the range for the memory size parameter was chosen based on memory model implementation and realistic values. The lower limit of three was chosen due to heuristics such as lexicographic that rely on the frequency of outcomes to determine decisions. To maintain the same parameter space for comparisons between memory models, this lower limit was applied to all models. The upper limit of ten is based on the findings of reasonable upper limits for memory capacity in Chapter 3: Modeling scale up of anthropogenic impacts from individual pollinator behavior to pollination systems.

## Scenarios

In the natural environment of bumblebees, the ability to quickly adapt to changes provides a clear advantage to individual pollinators. To improve our understanding of the underlying strategies and mechanisms of bumblebee decision-making, I simulate three sets of ecologically relevant scenarios that explore behavioral adaptation to change. In each scenario, individual simulations consist of one bee agent placed in an 80 x 80 patch virtual field with two species of plants uniformly distributed throughout. Each scenario also tests the eight different memory models found in Table 11 (scenario 1 does not test WSLS since each option always provides a reward). Simulated bees are initialized with full memories indicating that both plant choices are rewarding, in the same way that bees are trained in laboratory experiments to know that options are rewarding. For each experiment, BehaviorSearch was used for parameter estimation where the objective function was averaged over 100 model runs. Simulating the theoretical behavior of bees using different strategies can help identify key differences that can be used to inform laboratory experiments.

## Scenario 1: Response to Environmental Variation

Here we test bumblebee response to novel variation in their environment. This scenario simulates an ecological condition in which bees foraging in a static environment are exposed to novel variation in floral resources. In this scenario, bees initially forage in an environment where two options provide equal and constant reward. Bees then forage in a new environment where the same two options provide equally variable rewards. As discussed previously, a bee tested in an environment with one constant and one variable reward demonstrates suboptimal tracking, likely due to varying levels of risk-aversion[58]. This approach is designed to avoid simply testing risk aversion and can help identify how bees prioritize recent experiences when exposed to variation.

This experiment was designed to isolate any behavioral response to change in variability without changing the average rewards in the environment and executed by members of the Gegear lab. Bees were first individually trained on separate arrays, each with one of two colors of artificial flowers, orange or blue. Then, bees were allowed to forage on a constant environment array with both colors of artificial flowers where the amount of reward in each test tube was a constant 5ml of 25% sucrose solution. Bees were allowed to forage on the array until they returned to the hive. Next, the array was altered to be a variable environment such that half of each color of artificial flower contained 5µl of 10% sucrose and the other half of each color contained 5µl of 40% sucrose. Individual bees were then tested on the array until they returned to the hive. The color and reward of each visit by a bee was recorded. The first 80 visits of the training and testing foraging runs for eight individual bumblebees were used for consistency (some, but not all bees made visits beyond the first 80). See Table 13 for a summary of this experimental setup. The original experimental data included a total of three test environments recorded as separate foraging runs. The last two runs were not included for further analysis since the focus of this experiment is the difference in behavior of bees between the training and test environment.

## Scenario 2: Response to Variation in Reward Probability

In this scenario, we test bumblebee response to variation in reward probability. While scenario 1 tested response to variation between environments, here we examine variability and change in rewards. This experiment simulates an ecological condition where a change in resources requires bumblebees to reverse learned associations. Experiments in human psychology where reward probabilities are reversed have identified varying levels of stickiness – slowly adapting to change – and recency – quickly adapting to change in individuals[181,184,212]. Altering reward frequency can provide novel insight into the biases present in bumblebee decisions-making. Additionally, altering probabilistic reward schedules can also answer questions about the effects of partial reinforcement in bumblebees. Investigating bumblebee's response to variation in reward is important for furthering our understanding of bee decision-making and the role of memory in adapting to change.

### Scenario 3: Response to Variation in the Frequency of Change

Can bees adapt to rapid change, and how? When change is infrequent in a variable environment, memory provides an adaptive advantage in determining the most profitable decision. However, when change is extremely frequent a strong recency bias is crucial for an adaptive response[171,184]. This scenario tests bumblebee response to environmental change as the frequency of change varies. We explore the limits of memory in adaptive behavior by increasing the frequency of reward schedule reversals in reversal learning experiments. Modeling bumblebee decision-making is important for clarifying how individual bees treat early and later experiences when forming decisions. Elucidating how bees weigh recent memories is an important step in understanding how individuals can successfully adapt to changing environments.

# Results

## Response to Environmental Variation

We first consider how bumblebees respond in general to the variability of rewards in their environment. The data from this experiment (see Table 13) was summarized as the proportion of visits to orange flowers in blocks of twenty visits and averaged over the eight bees tested in the experiment.

*Table 13: Experimental scenario of changing variability in reward. The values refer to the percent sucrose of the 5μl solution in a test tube. Two species of plants are represented by the blue and orange coloring. Bees forage on the training environment for 80 visits, then forage in the test environment for an additional 80 visits.*

| **Probability of Reward** | | | |
|---|---|---|---|
| Training Environment | | Testing Environment | |
| Blue | Orange | Blue | Orange |
| 25% | 25% | 10% / 40% | 10% / 40% |

These results are plotted below in Figure 13. In the simulated version of this experiment, individual bee agents initially sampled the constant environment until their memory was full. They then foraged for 80 visits, at which point the simulation changed to the variable environment. After an additional 80 visits in the variable environment, the visitation choices and total reward collected were saved for analysis.

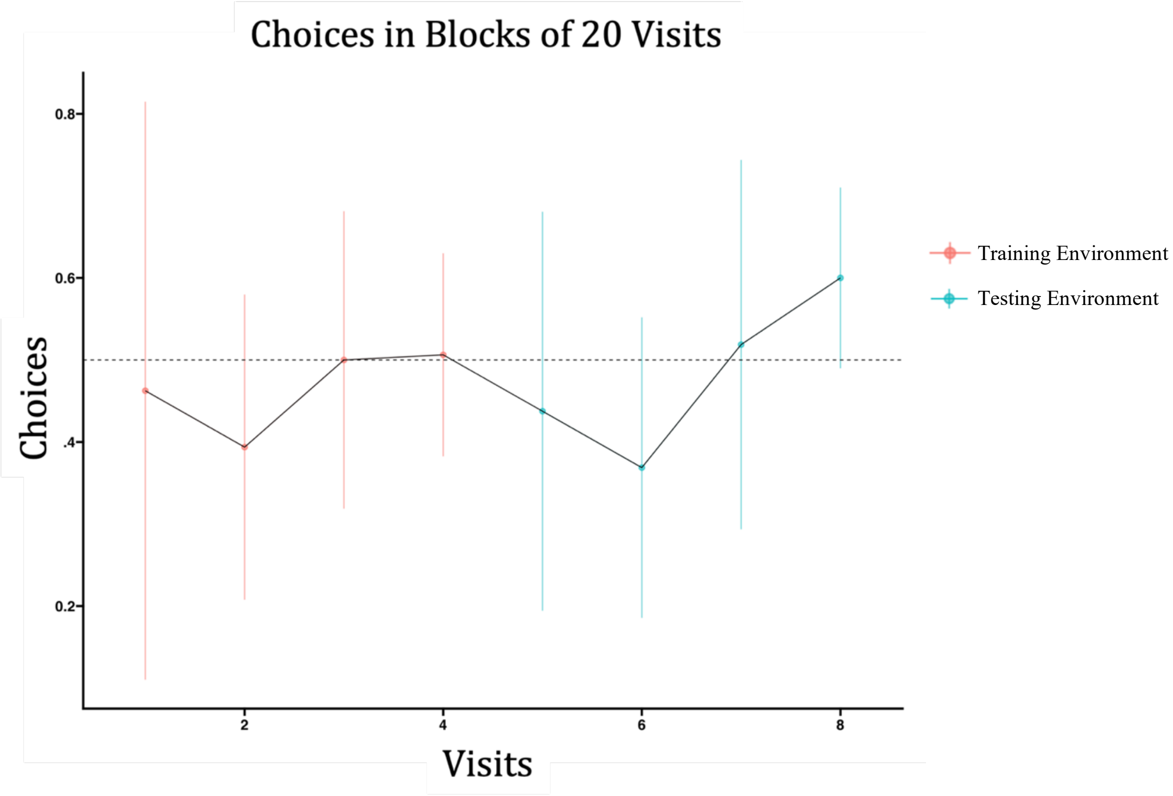

*Figure 13: Experimental data generated by Rob Gegear's lab. The plot shows visitation choices of 8 bees summarized into groups of 20 visits for a total of 80 visits per environment The x-axis shows blocks of visits and the y-axis shows the proportion of visits where the orange flower was chosen. Visits in the constant environment (reward of 25% sucrose) is shown in orange, and visits in the variable environment (10% or 40% sucrose) is show in blue). Each point represents the mean proportion of visits to orange flowers, with the bars representing the standard deviation.*

Although the experimental data demonstrates no clear statistical trend, the training and test environment seem to have different overall pattern in choice and variation. In particular, variation increases when bees switch environments indicating that there may be a response to change. However, given the limited sample size it is difficult to determine the true source of the observed variation. We proceeded to simulate the expected behavior of bumblebees with differing memory models to determine if the pattern of choice behavior matches the simulated bees. Parameter estimation was used to determine the optimal set of parameters that best matched the observed experimental behavior for each memory model.

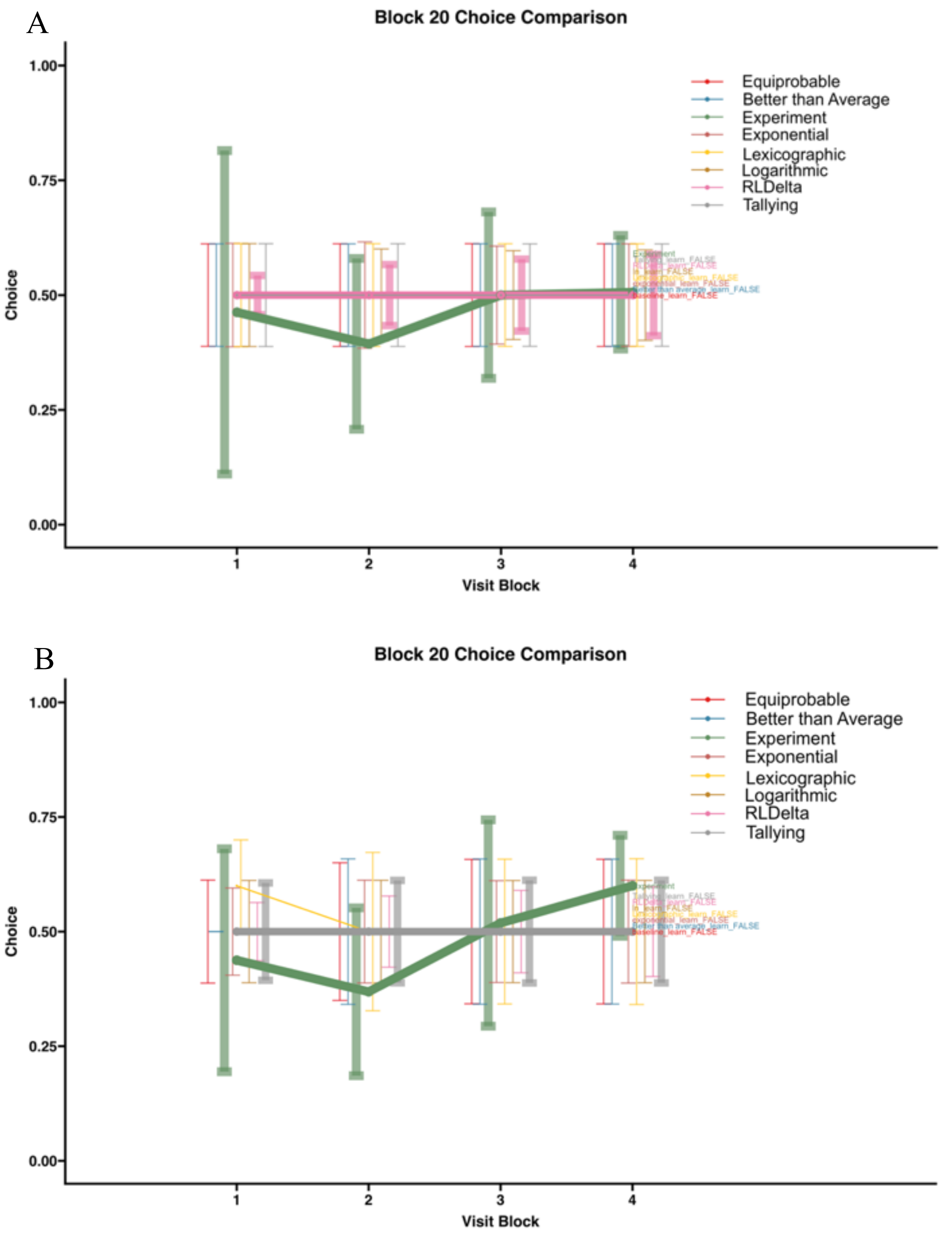

*Figure 14: A) Simulated bumblebees with different memory models fit to the experimental constant environment data in Figure 13. Experimental data is plotted in green. Error bars represent standard deviation. B) Simulated bumblebees with different memory models fit the experimental variable environment data in Figure 13. Fitting was accomplished using BehaviorSearch.*

Regardless of memory model, the simulations predicted random choices between the two available species of plants during the constant training scenario (Figure 14A). While we observed decreasing variation throughout the foraging runs of experimental bees, the simulated bees did not produce a similar change in variation.

Next, we considered the test environment where there is variation in reward but the same average reward. In the experimental bees, we observe some change in choice behavior over the foraging period. While we observed more variation in the simulated choice than in the virtual training environment, the overall pattern of random choice remains (Figure 14B). Every memory model produces similar behavior that does not match the observed experimental data.

### *Conclusions*

While the goal of this scenario was to determine how bumblebees respond to the introduction of variation in their environment and determine decision-making strategies, the experimental and simulated results are inconclusive. Each memory model tested in this scenario predicts similar choice behavior of equal visitation making it difficult to differentiate models. Additional experimental data is required to determine if bumblebees change their behavior in response to variation. In the future, a more complex experimental setup may be required to study the response to variation.

## Response to Variation in the Probability of Reward

The second scenario we consider is the effect of variation and change in reward probability on bumblebee behavior. This simulates an ecological condition in which floral resources are both variable, and rapidly change. This condition is modeled by extending the reversal learning paradigm beyond the traditional 100% chance-of-reward/0% chance-of-reward setup. Instead of having one option always be rewarding and the other option never rewarding, we test new reward schedules where options have an 80%/20% and 60%/40% probability of being rewarding. These alternative reward schedules make it more difficult to identify a change in the environment and therefore we hypothesize this environment should require a complex strategy to optimally exploit. Table 14 summarizes the three different simulation setups for this scenario. Additionally, this experimental setup has not been fully tested in a lab setting so we use the SimBee model to simulate the scenario, optimizing the nectar reward gained by bees using various memory models rather than fitting behavior to laboratory data. These simulations can help us make predictions and inform future experiments.

*Table 14: Simulated scenarios with changing reward probabilities. The values refer to the probability of bees receiving a reward – the reward stays the same throughout the scenarios. Two species of plants are represented by the blue and orange coloring. After 80 visits the reward schedules are reversed, and bees are allowed to forage for an additional 80 visits.*

| | **Probability of Reward** | | | |
|---|---|---|---|---|
| | Before reversal | | After reversal | |
| | Blue | Orange | Blue | Orange |
| **A** | 100% | 0% | 0% | 100% |
| **B** | 80% | 20% | 20% | 80% |
| **C** | 60% | 40% | 40% | 60% |

Each individual simulation starts with a bee agent with memory indicated that both species of plants are equally rewarding. The environment is initialized to have one plant species be the most likely to be rewarding, and the other species is the least likely to be rewarding. The environment does not change until the bees make 80 visits to plants, at which point the reward schedules reverse. Then the bees continue for another 80 visits. After another 80 visits, the simulation ends, and the fitness of the bee is measured as the total reward collected. Parameter estimation was used to determine the optimal set of parameters for each memory model. The results in Figure 15 show the optimal performance of bees using different memory models in this simulated scenario.

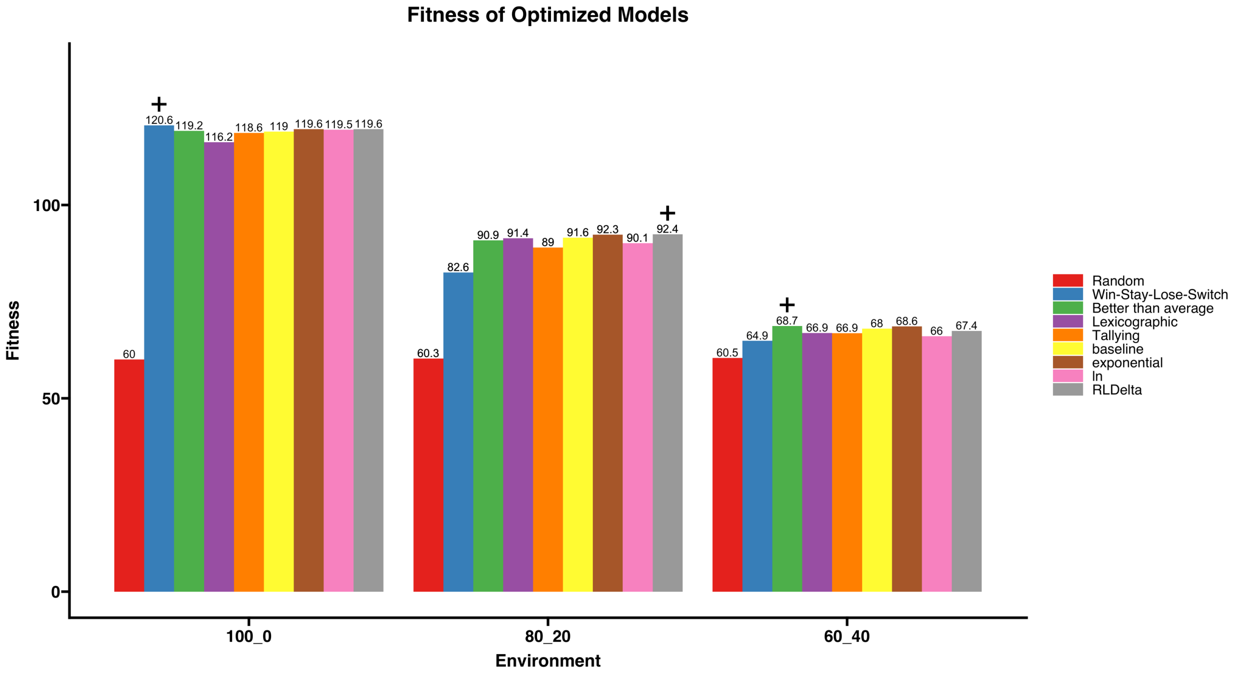

*Figure 15: Fitness of optimized models for simulated experiments. '+' indicate the model with the highest fitness (total reward collected).*

The 100%/0% reversal learning simulations demonstrates that near-optimal performance can be achieved with almost any memory model we tested. While the simple Win-Stay-Lose-Switch (WSLS) strategy performs the best, all the other optimized strategies except random choice collect near the maximum amount of reward from the environment. In the 80%/20% environment, WSLS actually performs worse than other strategies while reinforcement learning (RLDelta) achieves the best results. Since there is a difference in fitness between WSLS and other strategies, experimental approaches may be able to conclude if bees are using a memory-based strategy or not. Figure 16 provides a closer look at the 80%/20% environment and the number of correct choices the bee makes during the experiment summarized in blocks of 10 visits. The correct choices demonstrate a clear difference between the WSLS strategy and the other strategies tested in this scenario. In a traditional reversal learning experiment, the number of correct visits after a reversal would drop drastically before gradually increasing. In this case, using choices is a better method than overall fitness for identifying differences in behavior due to memory models.

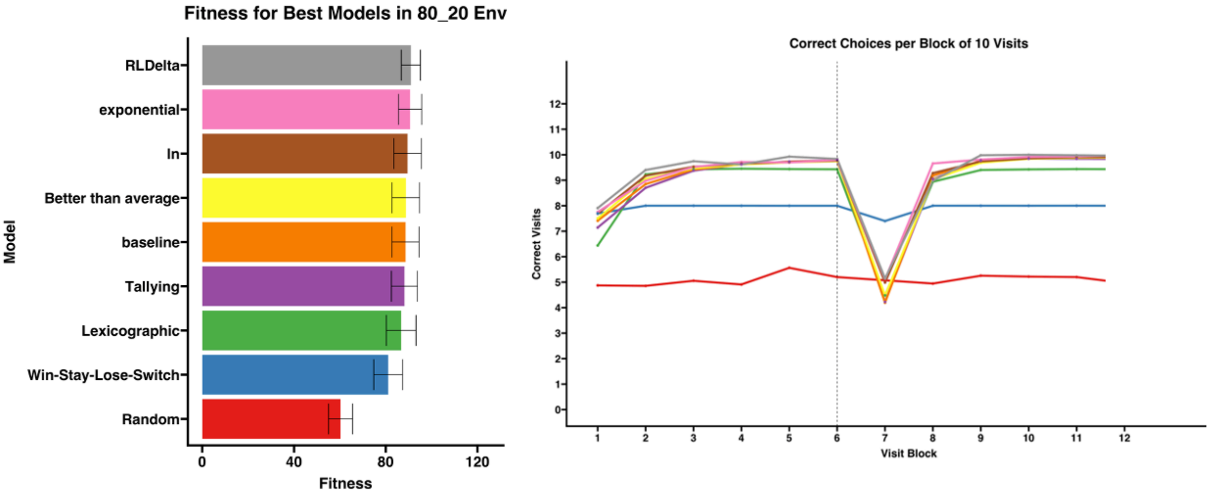

*Figure 16: Decision strategy performance with mean and standard deviation (left) and Correct Choices per Block of 10 Visits (right) for the 80%/20% environment. Fitness is measured as total reward collected. Dotted line shows the timing of the change in reward probabilities for the two species of plants*

Lastly, the 60%/40% environment shows relatively similar performance between the tested strategies. Although the heuristic Better than Average achieves the highest fitness, differentiating between strategies would be difficult in an experimental setting based on the theoretical performance in the simulated scenario.

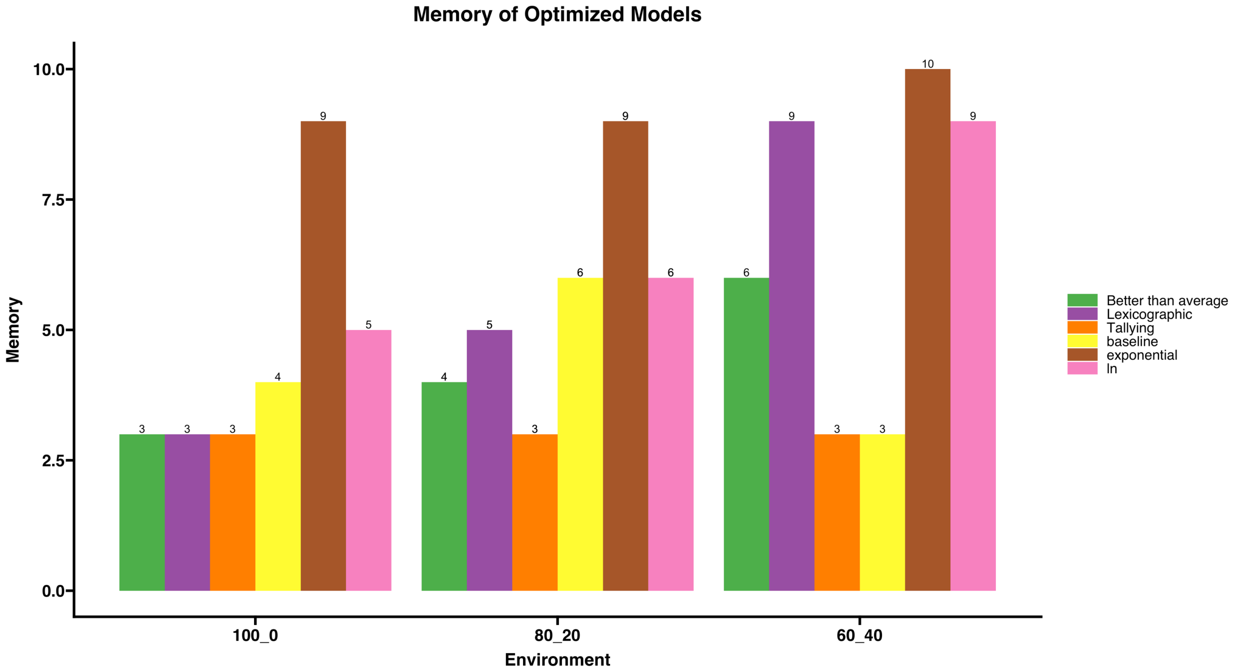

*Figure 17: Memory capacity of optimized models. Memory capacity is one of the parameters determined through parameter estimation as listed in Table 12.*

We also investigate the importance of memory capacity in the optimized strategies tested in this scenario (Figure 17). Across the three reward schedules, the exponential memory model required the largest memory capacity. Notably, the alpha values for the exponential model were also relatively low, indicating that both newer and older memories were being utilized (see Table 15). In contrast the Better than Average, Tallying, and Baseline strategies were optimized at much lower memory capacities indicating that a short memory works better than a long memory for these models. The Lexicographic and Logarithmic strategies increased memory capacity as environmental complexity increased. Since the memory capacity of bumblebees in the reversal learning paradigm is unknown, we cannot yet use these results to improve the accuracy of our simulated predictions; however, the parameterization of memory capacity can help determine which strategies are more likely to work with realistic memory limitations. In particular, the fact that Lexicographic and Better than Average strategies performs similarly to the more complex

Exponential and Logarithmic weighted memory models while using less memory suggests those strategies might be better representations of the actual bumblebee decision-making mechanism.

### *Conclusions*

The results of the simulations exploring the effect of variation in reward magnitude make it clear that more complex scenarios are needed to test memory and decision-making strategies in bumblebees. In the 100%/0% traditional reversal learning experiment, Win-Stay Lose-Switch is the best strategy and memory is not an adaptive trait. In contrast, the 80%/20% experiment may provide additional insight as memory is adaptive in that environment. Bees in laboratory experiments tend to follow patterns similar to the heuristics and memory models where they learn to visit the most rewarding option almost exclusively and take time to reverse preferences, and not patterns similar to those demonstrated by WSLS[159]. However, possible insight is limited since the memory-based strategies resulted in similar performance (results summarized in Table 15). An experiment focused on the number of visits a bee needs to reverse preferences with different memory models could better characterize the behavior of memory models. Overall, experiments studying bumblebee memory and decision-making need sufficiently complex environments to differentiate the performance of memory-based strategies.

*Table 15: Results for the top four decision-making strategies and memory models for the reward magnitude scenario. See Table 24 in Appendix S3 for the results of all models.*

| **fitness** | **memory size** | **memory model** | **memory alpha** | **softmax theta** | **env** |
|---|---|---|---|---|---|
| 120.6 | NA | Win-Stay-Lose-Switch | NA | NA | 100_0 |
| 119.64 | 9 | exponential | 0.7 | NA | 100_0 |
| 119.64 | 1 | RLDelta | 0.25 | 50.5547 | 100_0 |

| | | | | | |
|---|---|---|---|---|---|
| 119.2 | 3 | Better than average | NA | NA | 100_0 |
| 92.4 | 1 | RLDelta | 0.65 | 99.4151 | 80_20 |
| 92.34 | 9 | exponential | 0.25 | NA | 80_20 |
| 91.56 | 6 | baseline | NA | NA | 80_20 |
| 82.56 | NA | Win-Stay-Lose-Switch | NA | NA | 80_20 |
| 68.68 | 6 | Better than average | NA | NA | 60_40 |
| 68.62 | 10 | exponential | 0.05 | NA | 60_40 |
| 67.42 | 1 | RLDelta | 0.85 | 18.3484 | 60_40 |
| 64.86 | NA | Win-Stay-Lose-Switch | NA | NA | 60_40 |

## Response to Variation in the Frequency of Change

The third scenario we simulate with the SimBee model explores the effect of variation in the frequency of change in the environment on bumblebee decision-making. Since memory is not adaptive in the traditional 100%/0% reward schedule, we use the 80%/20% reward schedule from scenario 2 as the base for this scenario. Rather than change the magnitude of reward variation, we instead alter the number of visits between reversals. These experiments are summarized in Table 16.

*Table 16: Experimental design for scenario 3, frequency of change in the environment. The 80%/20% reward schedule remains consistent, but the number of visits between reversals decreases. We run each simulation for the same number of total visits, so bees in simulations with fewer visits between reversals experience more environmental changes.*

| | **Frequency of Reversal** | | |
|---|---|---|---|
| | Probability of Reward | | |
| | Blue | Orange | Visits between reversals |
| **A** | 80% | 20% | 60 |
| **B** | 80% | 20% | 40 |

| | | | |
|---|---|---|---|
| **C** | 80% | 20% | 20 |

Simulations are run for the same duration, 300 visits. Thus bees experience four reversals in experiment A, seven reversals in experiment B, and fourteen reversals in experiment C.

Similar to the previous scenarios, parameter estimation was used to determine the optimal set of parameters for each memory model in the given environment. The results in Figure 18 show the optimal performance of bees using different memory models in this simulated scenario. As expected in the 60-visit scenario, RLDelta is still the best performing strategy, much like the results of scenario 2. While the fitness of the models appears to increase with model complexity, the largest jump in fitness occurs between the WSLS and Lexicographic strategy. The other optimized memory models return somewhat similar fitness values. In the 40-visit environment where the frequency of environmental change is increased, the trend of fitness increasing with complexity is stronger. This indicates that the ability to place more decision-making weight on recent experiences, whether through recency bias by adjusting alpha or smaller memory capacity, provides an advantage in rapidly changing environments. Lastly, in the 20-visit environment the best performing strategy is the Logarithmic weighted memory model, but RLDelta and WSLS achieve a nearly similar fitness. The presence of WSLS in the top three strategies suggests that memory is not adaptive in such a rapidly changing environment.

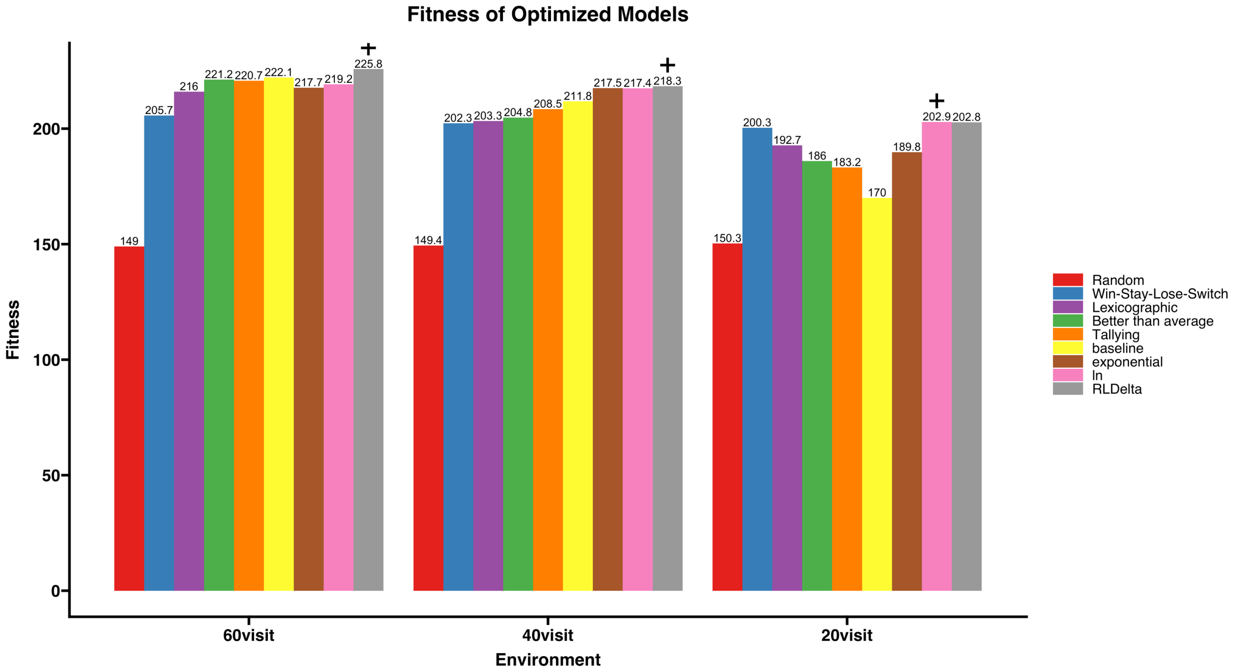

*Figure 18: Results for simulations on the frequency of change in the environment. '+' indicate the model with the highest fitness.*

For additional insight into the performance of the memory models, we investigate the memory capacity of the optimized models in this scenario (Figure 19). As the frequency of reversals increases, the optimized memory capacity decreases towards the minimum value of three across the board. These results strongly suggest that memory is a hindrance as the frequency of change increases, and bumblebees would likely avoid using complex decision-making strategies.

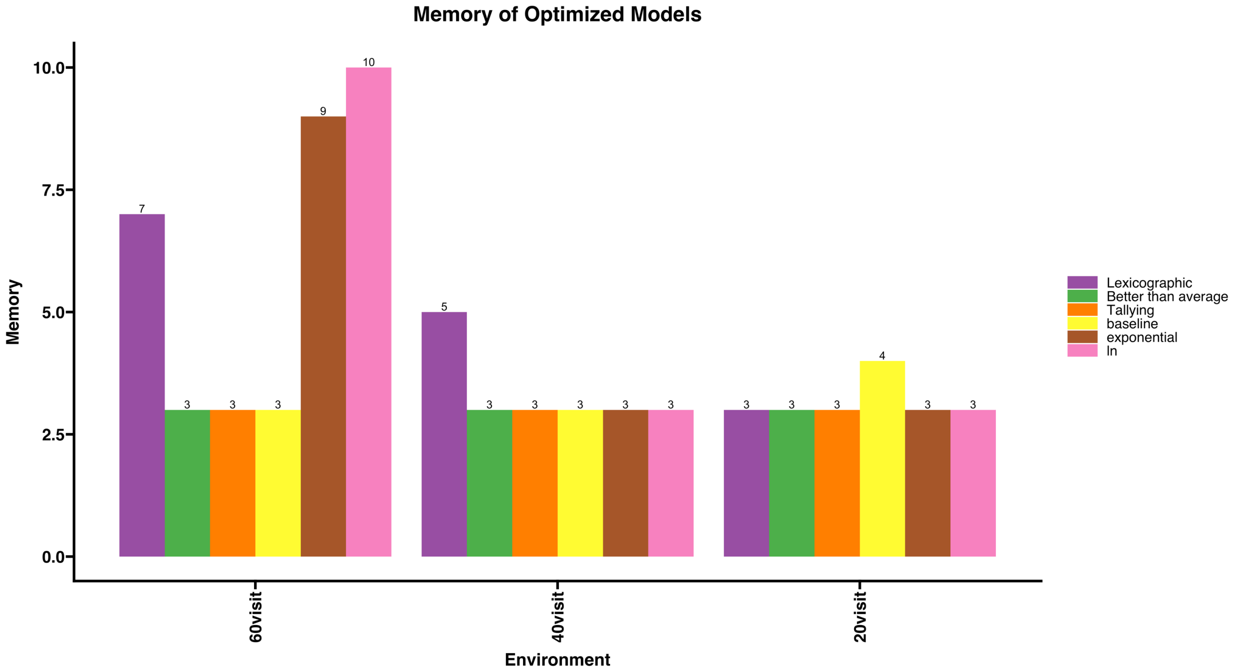

*Figure 19: Memory capacity of the optimized memory models in scenario 3. The value of three is the minimum memory capacity allowed; optimized models with a memory capacity of three indicate that additional memory capacity beyond three does not benefit bee fitness.*

### *Conclusions*

The analysis of simulation results for this scenario reveals two key trends. First, that increased memory model complexity generally resulted in better fitness when the environment rapidly changed in initial trials with 60 or 40 visits between changes. The RLDelta strategy was most often the optimal strategy, and the weighted memory models were the next best performing models in the 40-visit environment. Second, increased memory capacity does not seem adaptive when the environment changes extremely rapidly as in the 20-visit scenario. Instead, heavily weighting the most recent experiences correlated the strongest with high fitness in that environment. Importantly, simple strategies like WSLS can achieve such results just as well as RLDelta. Overall, the results suggest that experiments investigating the frequency of change highlight that memory is not always adaptive. The fitness and parameters of the optimized models are summarized in Table 17.

*Table 17: Results for the top four decision-making strategies and memory models for the frequency of change scenario. See Table 25 in Appendix S3 for the results of all models.*

| **Fitness** | **Memory size** | **Memory model** | **Memory alpha** | **Softmax theta** | **Visits / reversal** |
|---|---|---|---|---|---|
| 225.76 | 1 | RLDelta | 0.8 | 73.0256 | 60 visits |
| 222.12 | 3 | baseline | 0.5 | NA | 60 visits |
| 219.2 | 10 | ln | 0.65 | NA | 60 visits |
| 205.6667 | NA | Win-Stay-Lose-Switch | NA | NA | 60 visits |
| 218.32 | 1 | RLDelta | 0.85 | 25.6951 | 40 visits |
| 217.52 | 3 | exponential | 0.65 | NA | 40 visits |
| 211.84 | 3 | baseline | NA | NA | 40 visits |
| 202.3 | NA | Win-Stay-Lose-Switch | NA | NA | 40 visits |
| 202.92 | 3 | ln | 0.7 | NA | 20 visits |
| 202.76 | 1 | RLDelta | 0.95 | 46.5936 | 20 visits |
| 200.3333 | NA | Win-Stay-Lose-Switch | NA | NA | 20 visits |
| 192.72 | 3 | Lexicographic | NA | NA | 20 visits |

## Discussion & Future Work

This work uses agent-based modeling to improve our understanding of memory and decision-making in bumblebees by testing various strategies and generating predictions for expected behavior. The first scenario investigated response to a change in reward variability. Experimental results showed an increase in the variability of decisions made when bees were first introduced to the variable environment, indicating that bees might respond to change by visiting different options. However, the memory and decision-making models we tested using SimBee did not reproduce this increase in variation. It is possible that this response to change is not captured by recency bias and memory capacity parameters, and an additional mechanism is needed to properly model this response. Future directions could include identifying what mechanism underlies this response. Other approaches could include using different metrics, such as correct visits after reversals.

Furthermore, if in the future enough bumblebees were tested to quantitatively measure changes in the variation of choices, it could be possible to fit the simulated models to the variation in choices rather than the mean. Such an approach may help determine if the observed behavior fits any of the existing models of decision-making. Alternatively, since the primary subject of interest is the response to change, one could also model the change in choice behavior during the transition between environments. If one considers random choice for two blocks of visits as the accuracy criterion for learning the environment, one could use the number of visits until bees return to visiting both options equally as the time until recovery. The number of visits until recovery (equal rates of visitation to both options) would represent how quickly bees to adapt to the variable environment. This could then be used to design an objective function to fit simulated models to the experimental data. However, given the existing experimental data it is unclear whether the bees would reach the accuracy criterion in the variable environment.

Otherwise, one possible explanation for the inconclusive findings is that the bees did not treat the variable test environment as a continuation of the previous training environment and instead behaved as though they were entering a novel environment with no assumptions about the rewards of the available options. This could explain the increase in variation and could be tested by adding a third novel option to the test environment to determine if the bees are utilizing experiences from the training environment or sampling a new environment. If bees are using their memory to make decisions, we would expect them to continue visiting the options they know are rewarding, while bees that are sampling would make multiple visits to the novel option to assess its reward.

The second scenario explored the response to changing the probability of reward in a reversal learning experiment. Bumblebees have successfully reversed preferences in reversal learning experiments with one option always rewarding and one option never rewarding, but our modification of the reward probabilities in reversal learning has not been tested in bumblebees. The results of our simulations demonstrate that the 100%/0% reward schedule in reversal learning is insufficient to distinguish between simple rule-based strategies like WSLS and more complex memory models. This finding reflects the difficulty of drawing conclusions about decision-making and memory from similar reversal learning experiments. Since it is likely that many different models would fit existing reversal learning experimental data that use the same reward schedule, it would be challenging to narrow down which decision-making strategies and associated memory capacities are the best fits. Increasing the complexity of the environment to an 80%/20% reward schedule reveals differences in fitness between theoretical decision-making strategies but differences in choice behavior between WSLS and other models is clearer. Here, the differences in decision-making were clear in the choices made immediately after the reversal. Weighted memory models and reinforcement learning can reproduce the effects of recency bias in bumblebee behaviors in ways that are not captured by simpler models since memory is an adaptive trait in the 80%/20% scenario. Testing the 80%/20% scenario in a laboratory would help validate these observed differences and provide data for determining a range of potential values for memory capacity. Future modeling work could also explore the sensitivity of the memory and recency parameters in more depth, as the BehaviorSearch parameter optimization does not explore suboptimal parameter value combinations in a robust way. However, the models tested in this scenario all predict that bees learn to reverse their preferences relatively quickly which is something that has not been observed in similar experiments with rats[185,186]. Specifically, experiments testing the effect reward probability on rat learning found that rats only reversed their initial learned preference when reward was guaranteed (100%) and otherwise failed to learn the reversal. Additional laboratory experiments would be helpful to determine if these predictions about partial reinforcement are also observed in bumblebees, or if bumblebees exhibit radically different behavior. If bees do not behave as predicted by our simulations, more complex models may be needed to accurately reproduce bumblebee responses to variation in reward probabilities.

Future directions for expanding on scenario two should include identifying experiments where differences in memory and decision-making produce markedly different behavior. Drawing inspiration from human psychology, experiments where rewards change gradually over time may provide the desired complexity. Konstantinidis et al. (2020)[181] designed an experiment where the probability of reward either increased or decreased over time and fit participant choices to several learning models including reinforcement learning. Doing so helped the researchers identify ‘adaptive’ and ‘non-adaptive’ individuals that exhibited different behaviors, as well and identifying the importance of recency in adaption to change. Testing bees in an environment where probability of reward for the best option decreases over time would provide useful insight into how strongly bees are influenced by initial experiences and could provide the complexity needed to separate decision-making strategies. Alternative ways of adding complexity to experiment testing decisions from experience include non-monotonic change – periodic change for example – and increasing the number of options to choose from. However, the increasing complexity of these scenarios does present challenges for experimental implementation. Overall, there are many possible variations for future experiments where simulations could provide useful predictions before laboratory tests are attempted.

In scenario three, we consider how bumblebees respond to increasingly frequent environmental change. The results of our simulation confirm that memory is key to optimally exploiting an environment with infrequent changes while memory is not adaptive for exploiting rapidly changing environments. How can bees adapt to both frequent and infrequent environmental change? Models such RLDelta and the logarithmic weighted memory model are flexible enough that bees can achieve high levels of performance in both the 60-visit serial reversal environment and the 20-visit serial reversal environment, indicating that recency bias is a possible mechanism for explaining how bumblebees can adapt to changing environments. The ability of bumblebees to alter their recency bias to fit their environment and the similarity of weighted memory models could explain why bumblebee behavior appears to match many different models and strategies in the literature.

Finally, the theoretical models implemented and tested in chapter four are only a small subset of possible models. Variations on Win-Stay Lose-Switch such as ‘deliberative-decisive’[190] where several unrewarding visits are needed to trigger a switch in preference may also serve as useful comparisons along with other heuristics[201] used in human psychology. Other versions of reinforcement learning are also used in human psychology, such as the decay learning rule where unchosen options are also discounted[181]. The current SimBee model does not have a decay parameter option, but future extensions could implement this. Furthermore, other models such as Bayesian decision models, optimal diet models, and contextual models (where individuals change strategies based on environmental context clues) could be added as comparisons for a more comprehensive study of decision-making in bumblebees.

# Chapter 5: Deep Learning Behavioral Classification of Bumblebee Videos

## Introduction

As digital cameras and video content has become ubiquitous in human society, the need for accurate behavior classification in videos has driven rapid advances in research on video action recognition of human behaviors. Behavior recognition consists of identifying some behavior from a series of observations, a task that has traditionally been very challenging. While the classification of actions from a sequence of image frames seems like an extension of image classification, progress in video behavior analysis has been slower due to several complicating factors. Video analysis requires not only spatial information but also temporal context across multiple frames. The consequence of retaining this information for videos of substantial length is huge computational costs[213,214]. Additionally, the process of collecting, annotating, and storing videos is more difficult than for single images, resulting in fewer and smaller datasets.

Initial approaches used specially designed features and frameworks to capture specific human-based actions[215]. Augmented by markers or tags, these early models could identify human motions given controlled conditions but failed to generalize well with unscripted videos. Additionally, these manual approaches were time-consuming to train on the ever-increasing amount of video data available. Advances in machine learning over the past decade have resulted in an explosion of network architectures and applications. Different techniques such as optical flow, Hidden Markov Models (HMM), and more recently deep learning have been applied to behavior recognition with varying success[215–222].

Despite the inherent difficulty of the task, there have been several novel approaches to the problem. In general, these methods fall into two main groups: methods using hand-crafted motion features and deep learning methods[213]. In hand-crafted methods, some interesting features are obtained from the raw pixels of the video frames and used to perform the recognition. This approach offers several advantages, namely that the features of the model are explicitly known, and customized features require relatively little data to train. However, hand-crafted features are usually not robust and often fail when used on data with more variability than the training set[215].

In contrast deep learning methods learn features through their network, reducing the difficulty of feature extraction but increasing the data requirements. Due to this, deep learning methods tend to work better on complex problems and can be more easily generalized to new problems as long as data is available. As large video datasets have become available, deep learning methods have become the recent focus of behavior recognition research.

## Video Action/Behavior Datasets

There are many large-scale annotated datasets for human video action recognition. Most consist of videos obtained from YouTube or movies and are created by breaking videos into clips and using resources such as Amazon's Mechanical Turk to add appropriate class labels[223]. Some datasets such as HMDB51[224] and UCF101[225] have a single label per short clip, while others like ActivityNet[226] and AVA[227] have multiple actions per video, with annotated start and end times. Datasets with temporal annotations are designed to enable temporal action recognition within videos, and some datasets take this further by adding boundary boxes within videos for spatio-temporal action recognition. However, adding spatio-temporal annotations greatly increases the complexity and cost of creating such datasets, making the existing ones extremely valuable. Although the dataset Moments In Time (MIT) has single-class non-temporal annotations, it is notable for intentionally including animals, objects, and natural phenomena in addition to people[228]. These human-oriented action recognition datasets are summarized in Table 18. While datasets for other organisms such as mice and flies exist, they are rarely designed for public use.

*Table 18: Summary of public datasets for human action recognition.*

| DATASET | YEAR | # OF SAMPLES | ACTION CLASSES | LABEL TYPE (# LABELS PER CLIP) | TEMPORAL ANNOTATIONS |
|---|---|---|---|---|---|
| **HMDB51[224]** | 2011 | 7K | 51 | Single | No |
| **UCF101[225]** | 2012 | 13.3K | 101 | Single | No |
| **ACTIVITYNET[226]** | 2015 | 28K | 200 | Multi | Yes |
| **YOUTUBE8M[229]** | 2016 | 8M | 3862 | Multi | Yes |
| **AVA[227]** | 2017 | 385K | 80 | Multi | Yes |
| **KINETICS400** | 2017 | 306K | 400 | Single | No |
| **MIT[228]** | 2018 | 1M | 339 | Single | No |

Intuitively, video behavior understanding needs information about motion. This is achieved using optical flow, the relative motion of objects, surfaces, and edges within a visual scene. Optical flow provides more information than the original RGB image because it combines frames to provide orthogonal information. In doing so, it effectively removes static background and displays movement information, creating a simpler learning problem (See Figure 20).

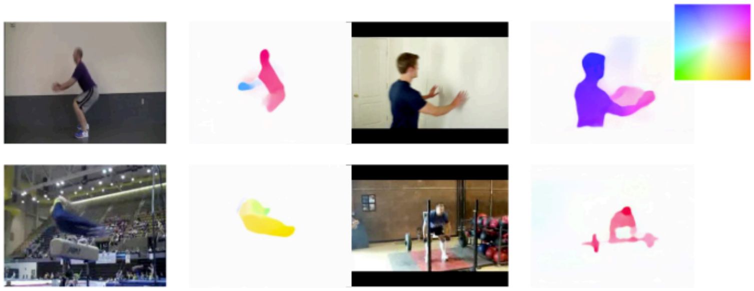

*Figure 20: Visualization of optical flow. Left is the original RGB image and right is the estimated optical flow. Color of optical flow indicates the directions of motion, and the color-coding scheme is shown in top right. Taken from Zhu et al. 2020[223].*

Two-stream networks extend the use of optical flow by using RGB data as input for the spatial network (or stream) to learn visual appearance, and optical flow data as input for the temporal network to learn motion information. While this approached initially worked well with very short clips, initial two-stream networks had difficulties capturing longer temporal information[230]. Temporal Segment Networks (TSN) was the first to combine the two-stream architecture with video sampling to achieve video-level behavior recognition. It achieves this by dividing videos into several segments uniformly distributed along the temporal dimension. It then samples from each segment and forwards them through the network. An aggregated consensus is determined from the output of each segment for each stream, which is then fused for a final consensus. TSN achieved a then record accuracy of 94% on the UCF101 dataset and inspired many other models (See Figure 21)[223]. TSN was also used for video behavior classification of pigs and achieved 98.99% accuracy[231].

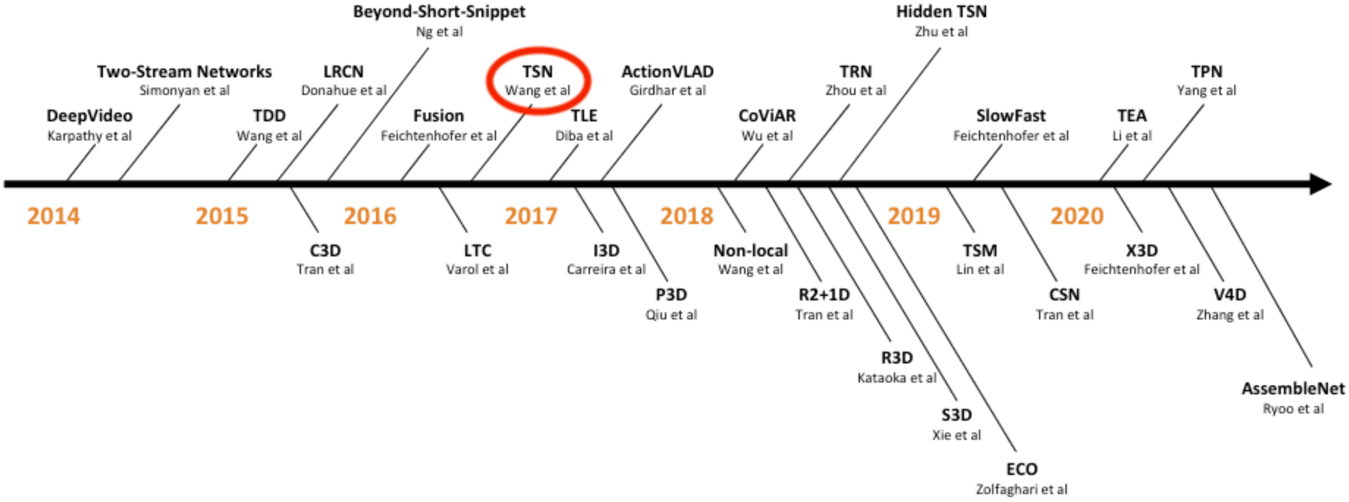

*Figure 21: A chronological overview of recent representative work in video action recognition. TSN, circled in red, is the network used in this work. Image taken from Zhu et al. 2020[223].*

Since TSN, there have been several noteworthy deep learning models that have advanced the field of video behavior recognition. I3D[232] was the first model to bypass the computationally intensive step of pre-computing optical flow using 3D convolutional neural networks by pre-initializing the weights for optical flow from older models. Temporal Shift Module (TSM)[233] performs temporal modeling without 3D convolution or optical flow by shifting parts of its channels along the temporal dimensions, thus achieving temporal information without optical flow. VideoLSTM[234] (not shown in Figure 21) takes this one step farther by using a specific type of recurrent neural network called LSTM to achieve the same temporal information as TSM with additional attention maps for action localization. VideoLSTM can both classify and localize actions from only an action class label.

While the state of deep learning video behavior classification models has thus made numerous advances since TSN, these advances have in part been enabled by increasingly large and complex datasets. Additionally, while the implementation of TSN has been redesigned and documented for public use through the MMAction2 Toolbox[235], more recent models such as TSM and videoLSTM have not received the same treatment so far. Considering the ease-of-implementation, lack of large-scale bumblebee datasets, and proven track-record on non-human organisms, we focus our efforts on using TSN for video behavior classification.

## Non-human Behavior

While deep learning methods have recently begun to achieve high levels of accuracy (>90%) in classifying human behavior, the same is not generally true for non-human behaviors. Unlike action recognition in humans where an action such as "running" looks similar for most individuals, animal behaviors are more specialized. Consequently, methods using hand-crafted features can work well for the organism they are designed and trained on – especially in controlled environments when the subject is always at the same angle – but do not generalize or transfer to other species. Experiments testing the generalizability of I3D (a deep learning method proposed shortly after TSN) demonstrate this issue; while I3D reached an accuracy of 80.9% on

HMDB-51 and 98% on UCF-101, it only reached 36% accuracy when tested on a dataset with thirty-two animal categories[236]. It is clear that animal behavior is a complex topic that must be approached one species at a time for the best results.

Although advances with deep learning models do not seem to transfer generally to different animals, researcher have seen far greater success when focusing on a single species. A study using TSN to detect pig behaviors found that the network was able to recognize pig feeding, lying, walking, scratching, and mounting behaviors with an average of 98.99% accuracy. A similar study on fox behaviors achieved a mean average precision of 99.91% when classifying the behaviors of multiple foxes[237]. However, these studies used datasets created by filming larger animals in closed environments with controlled lighting, ensuring low variance in video quality. Furthermore, the deep learning models in these studies could successfully localize individuals, something that does not seem to always hold true when studying smaller insects. While approaches for classifying insect behavior exist, they generally rely on markers for identification[44,238–240]. This works well in laboratory studies where markers can be placed on insects but is less useful for studying the behavior of wild organisms. It seems likely that deep learning methods can achieve some level of success in classifying bumblebee behaviors without markers and hand-crafted features, but deep learning models such as TSN have not been tested on real-world datasets, in part because generating such datasets is a time-consuming process.

## Bumblebee Behaviors

For eusocial insects such as bumblebees, the gathering of nectar and pollen is critical for the growth and maintenance of the colony[241]. Nectar is the primary source of energy for the queen and the workers, while pollen supplies protein for developing larvae[68,242,243]. Since the collection of nectar and pollen varies temporally for bumblebees in a colony, knowledge of the behaviors exhibited by bumblebees can provide important insights into the health and lifecycle of wild colonies without physically examining them[65,111]. Specifically, newly established colonies grow by producing new cohorts of workers while established colonies switch to only making males and new queens. Once a colony switches, it ceases to produce new workers[244]. Tracking the behavior of foraging bumblebees can provide indicators of this shift before the demographics of the hive change[120]. Behavior information can also provide insight into task specialization in colonies[245–247]. Overall, behavioral information at all levels can help improve our understanding of bumblebee individuals and colonies.

Although there are many different behaviors to consider when studying bumblebees, we focus on five behaviors that are likely to be present and identifiable in videos submitted by the citizen scientists. Additionally, the number of behavioral classes is also limited by the ease of identification. The behaviors included in the dataset for this work are nectar foraging, pollen foraging, flight, approaching a flower, and departing a flower.

## Nectar Foraging

Nectar foraging is one of the primary foraging behaviors of bumblebees. For most foragers, this action involves extending the probiscis deep into a flower to check for the presence of nectar (see Video 1). If nectar is present, the forager my spend an extended amount of time in the same position while it extracts all available nectar from the flower. This behavior of keeping the head still as the forager probes a flower for nectar contrasts with pollen foraging, where the bee does not remain still. Some bumblebee species use an alternate behavior known as nectar robbing, in which the bee feeds on nectar by biting a hole in the base of the flower[248]. However, these bees are not present in the dataset used here.

*Video 1: Clip of a Bombus impatiens female gathering nectar from a flower.*

## Pollen Foraging

Pollen gathering is an important activity for foragers, as pollen provides the protein necessary for a growing colony. Bumblebees collect pollen by shaking it loose from flowers and letting it stick to their furry bodies[59]. They then groom themselves and wipe the pollen onto stiff hairs on their hind legs or abdomen, often referred to as pollen baskets. Once a forager's pollen baskets are full, the forager returns to the hive and deposits the pollen. This behavior can be identified by a forager's rapid movements meant to dislodge pollen from the flower and is shown in Video 2 below.

*Video 2: Clip of a Bombus impatiens female gathering pollen from a flower. The pollen baskets can be observed on the bee's hind legs.*

## Flight

Flight is a behavior of foraging bumblebees that correlates with pollination activity[249]. The ability to track bees as they move from flower to flower is useful in behavioral contexts, and most extended videos of bumblebees include flight (see Video 3). Since flight behavior is often present in submitted videos and is distinctly different from nectar and pollen foraging, behavioral classification of flight serves as an important step for deep learning models.

*Video 3: Clip of a Bombus vagans female flying between flowers.*

## Approach

*The behavior of a bumblebee approaching a flower (from flight) has a strong temporal component (see* Video 4). ‘Approach’, and its corollary ‘departure’ both involve a transition from one behavior – nectar/pollen foraging – to another – flight. Additionally, the timeframe of this behavior is much shorter than other behaviors. Correctly classifying such behaviors can help provide temporal information that improves the accuracy of other behavior classifications.

*Video 4: Clip of a Bombus vagans female approaching a flower.*

## Departure

Similar to ‘approach’, the ‘departure’ behavior always occurs after foraging and before flight. Correctly classifying this behavior can help provide temporal information that improves the accuracy of other behavior classifications. Video 5 shows an example of a bumblebee departing from a flower.

*Video 5:Clip of a Bombus impatiens female departing from a flower.*

## Using Citizen Science Data

It is important to address the challenge of accelerating the validation of citizen-generated data by utilizing modern deep-learning methods to automate behavior classification of bumblebee videos. While the manual analysis of videos is impractical, recent advances in video behavior classification of human behaviors have demonstrated incredible success in accurately classifying a wide range of behaviors from amateur videos[213,225,228,250]. Specifically, multi-steam LSTMs have achieved over 95% accuracy on large video datasets such as UCF101[225], HMDB51[224], ActivityNet[226], and others. However, it is unclear how well such deep-learning architectures and networks transfer to bumblebees, in part due to the lack of annotated video data for bumblebees. For this research, we use the two-stream convolutional neural network architecture proposed as TSN[251] in 2016 because the authors have developed a toolbox[235] for applying their work to new datasets and these tools have been successfully used in automated video behavior recognition of pigs[231]. Using citizen science videos of bumblebees to create a novel behavior recognition dataset, we retrain TSN for bumblebee behavior classification and present the preliminary results below.

## Beecology Dataset

The data for this work comes from the Beecology Project, a collaboration to develop and maintain an accessible method for citizen scientists to contribute field observations of native pollinators. At the time of inception, this dataset included 434 user-submitted videos. The majority of these – 374 videos to be precise – record nectar gathering behavior, while the remaining 60 are pollen gathering. Most bumblebees featured in these videos were *impatiens* (181) and *vagans* (124). Additionally, most of these bumblebees were females/workers (331), with 92 males, and two queens recorded. Across the dataset, bumblebees visited a total of 88 different plant species. See Table 19 for the specific parameters of these videos.

*Table 19: Specific parameters of the bumblebee behavior dataset.*

| **Item** | **Parameter** |
|---|---|
| Behavior class | 5 |
| Behavior name | Nectar, Pollen, Departure, Flight, Approach |
| Resolution | 320 x 240 |
| Frame rate | 25 fps |
| Video Mean Duration | 7 seconds |
| Storage format | AVI |

## Methods

### Data Processing

From the original 434 videos, 237 were selected for further annotation. Each video was converted into frames and manually split into different behaviors. These clips were labeled in the format v_**X**_g**Y**_c**Z**_s**A**_d**B**_f**C**_h**D**.avi where **X**=Behavior, **Y**=group #, **Z**=clip #, **A**=bee species, **B**=bee gender, **C**=flower genus, **D**=flower species. Repeated behaviors from the same original video were identified by incrementing the clip number.

The 237 processed videos were then divided into a training set and test set in a 70%/30% split. The training set was further divided by a 70%/30% split into a new training set and validation set for an overall 50/20/30 training-validation-testing split (see Table 20).

*Table 20: Summary statistics of the 50%/20%/30% training-validation-testing split of the 237 processed videos.*

Training

| Behavior | Count |
|---|---|
| Pollen | 17 |
| Nectar | 41 |
| Flight | 19 |
| Departure | 27 |
| Approach | 13 |
| Total | 117 |

Validation

| Behavior | Count |
|---|---|
| Pollen | 7 |
| Nectar | 18 |
| Flight | 8 |
| Departure | 11 |
| Approach | 5 |
| Total | 49 |

Testing

| Behavior | Count |
|---|---|
| Pollen | 10 |
| Nectar | 25 |
| Flight | 11 |
| Departure | 17 |
| Approach | 8 |
| Total | 71 |

The motion information in the videos is very important for behavior recognition, and the optical flow can incorporate motion information such as the direction and speed of a moving target. The OpenCV implementation of the TVL1 optical flow algorithm was used to generate optical flow values in the horizontal and vertical directions for all clips using the same process as the original TSN model[252]. The RGB images were used as the input of the spatial convolution network (first stream) and the optical flow was used as the input of the temporal convolution network (second stream). The results of the two streams are fused together to obtain the video behavior classification.

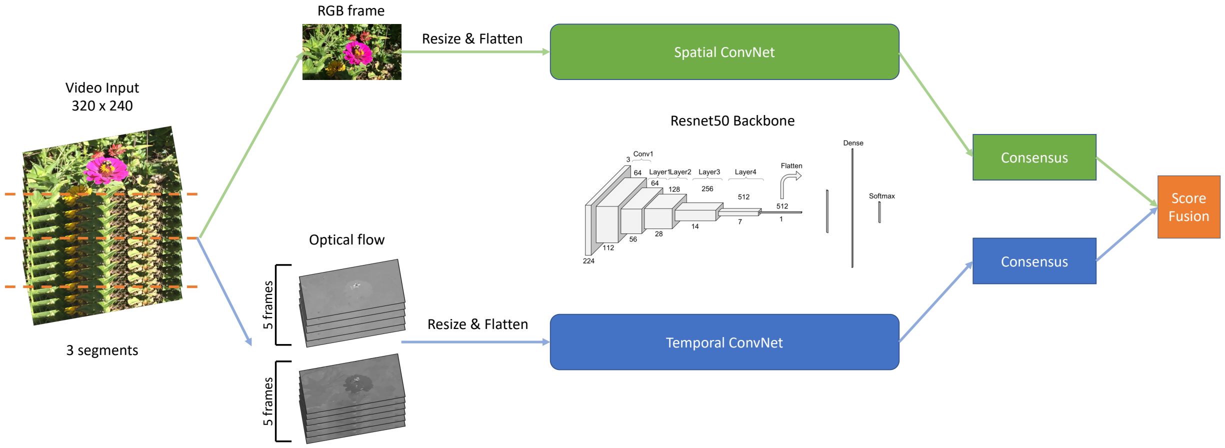

*Figure 22: Temporal segment network: One input clip is divided into three segments and a short snippet is randomly selected from each segment. The class scores of different snippets are fused by a segmental consensus function to yield segmental consensus, which is a video-level prediction. Predictions from both modalities are then fused to produce the final prediction. ConvNets on all snippets share parameters. Architecture based off of TSNs: Towards Good Practices for Deep Action Recognition*[251]*. ResNet50 diagram taken from "Understanding and visualizing ResNets"*[253]*.*

## Baseline Model

To better characterize the Beecology dataset, the 'most frequent' classification strategy is used to generate a baseline for more complex models. For this dataset, the 'most frequent' strategy always classifies a clip as nectar foraging. When fit and tested using the data splits shown before, this produces an accuracy of 35.2% (see Table 21). Based on these initial tests, more complex methods should produce higher levels of accuracy. If the level of accuracy remains similar to the baseline strategy, that method may not be capable of classifying the existing Beecology data.

*Table 21: Classification statistics for the 'most frequent' baseline strategy. The baseline accuracy is 0.352.*

| | PRECISION | RECALL | F1-SCORE | SUPPORT |
|---|---|---|---|---|
| **NECTAR** | 0.35211268 | 1 | 0.52083333 | 25 |
| **POLLEN** | 0 | 0 | 0 | 10 |
| **APPROACH** | 0 | 0 | 0 | 8 |
| **DEPARTURE** | 0 | 0 | 0 | 17 |
| **FLIGHT** | 0 | 0 | 0 | 11 |
| **ACCURACY** | 0.35211268 | 0.35211268 | 0.35211268 | 0.35211268 |
| **MACRO AVG** | 0.07042254 | 0.2 | 0.10416667 | 71 |
| **WEIGHTED AVG** | 0.12398334 | 0.35211268 | 0.18339202 | 71 |

## Experimental Parameter Settings

The network weights for the TSN model were initialized using the pre-trained ResNet50 network from ImageNet[254]. The learning and momentum policies were varied during experiments to determine optimal parameters. The optimization method was stochastic gradient descent. For both networks, the dropout ratio was 0.4 and total epochs were 75. During training, the number of snippets (segments taken from a single clip, refer to Figure 22) was 3.

## Experimental Environment

The model experiments were all conducted on the Worcester Polytechnic Institute (WPI) Turing server. Each experiment was run using either a V100 or A100 NVIDIA GPU depending on availability. The server runs Red Hat Enterprise Linux 7.3, and the core software included

CUDA 11.1.0, cuDNN 8.1.1, OpenCV 4.2.0, and pytorch 1.9.0. Denseflow was used to generate the optical flow of videos[255].

## Preliminary Results

The training dataset was used as input into the TSN networks while the learning rate and momentum policies were varied. The top 1 accuracy - the proportion of clips for which the predicted action class matches the true action class - and mean class accuracy – the average accuracy of the predicted action class for all action classes - were evaluated using the validation dataset every five epochs. After training was complete, the best model was evaluated on the test dataset. For RGB spatial network, the best test top 1 accuracy achieved was 38.77%. The optical flow temporal network achieved a much higher accuracy on the test set at 57.75% top 1 accuracy. These experiments are summarized in Table 22 for the RBG network and Table 23 for the optical flow network. Considering the effect changing the learning rate, it appears that a relatively small learning rate is beneficial for the initial training of the model. Overall it seems like most learning occurs in the early training iterations, so an initially small learning rate provides better results than altering momentum policies.

*Table 22: Model performance for the RGB network during training and testing.*

| | | | | | | VALIDATION | | | TESTING | |
|---|---|---|---|---|---|---|---|---|---|---|
| **VERSION** | **Learning Rate** | **Momentum** | **Weight decay** | **Learning rate** | **Momentum policy** | **Best Epoch** | **Top1acc** | **Mean cls acc** | **Top1acc** | **Mean cls acc** |
| **1** | 0.00128 | 0.9 | 0.0005 | 0.00128 | NA | 70 | 0.26531 | 0.18968 | 0.2254 | |
| **2** | 0.00028 | 0.9 | 0.0005 | 0.00128 | NA | 15 | 0.38776 | 0.22525 | 0.3662 | 0.2118 |
| **3** | 0.005 | 0.9 | 0.0001 | Cyclic learning policy | Cyclic Momentum Policy | 50 | 0.38776 | 0.26121 | 0.2817 | 0.1713 |
| **4** | 0.005 | 0.9 | 0.0001 | Cyclic learning policy | Cyclic Momentum Policy | 10 | 0.36735 | 0.2 | 0.3673 | 0.2 |

*Table 23: Model performance for the optical flow network during training and testing.*

| | | | | | | VALIDATION | | | TESTING | |
|---|---|---|---|---|---|---|---|---|---|---|
| **VERSION** | **Learning Rate** | **Momentum** | **Weight decay** | **Lr policy** | **Momentum policy** | **Best Epoch** | **Top1acc** | **Mean cls acc** | **Top1acc** | **Mean cls acc** |
| **1** | 0.005 | 0.9 | 0.0001 | Step learning policy | NA | 5 | 0.4898 | 0.35919 | 0.5634 | 0.42 |
| **2** | 0.001 | 0.9 | 0.0001 | Step learning policy | NA | 20 | 0.57143 | 0.4303 | 0.5775 | 0.44 |
| **3** | 0.001 | 0.9 | 0.0001 | Cyclic learning policy | Cyclic Momentum Policy | 9 | 0.5102 | 0.2826 | 0.5352 | 0.4118 |

The change curve of the loss functions and accuracy rate of the TSN networks helps visualize the results of training over time. For the RGB network, the loss (a representation of the error between the predicted and actual action classes for training data) rapidly decreases at the start of training before quickly leveling off (Figure 23), a behavior that can be associated with a well-fitting model. However, the overall low accuracy (Figure 24) indicates that the RGB network is unable to properly predict the correct class given the dataset. The loss curve for the Optical Flow network (Figure 23) on the other hand does not demonstrate that the model is fitting the data, although it achieves an overall higher accuracy than the RGB network (Figure 24). Both accuracy curves reach their peak early in training, indicating that additional training is not helpful for improving accuracy.

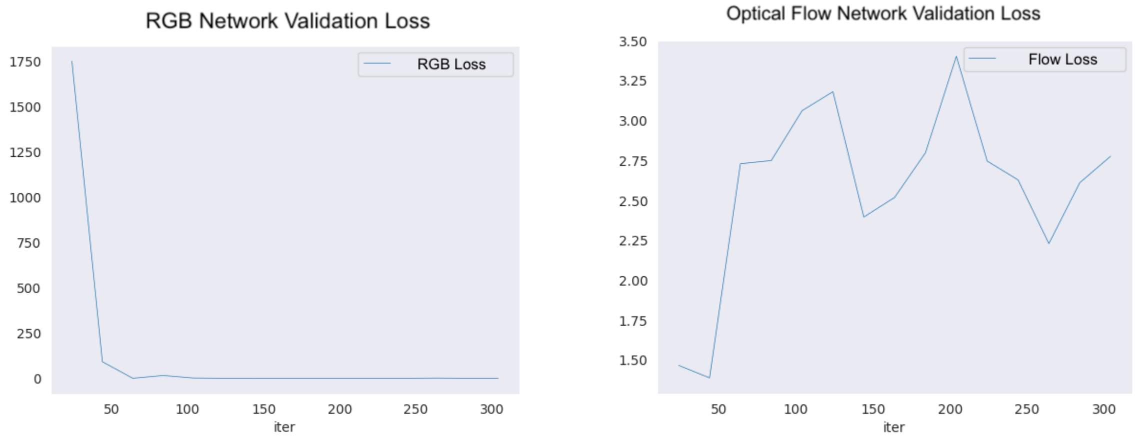

*Figure 23: RGB and Optical Flow network loss function during training. The loss function is cross entropy loss.*

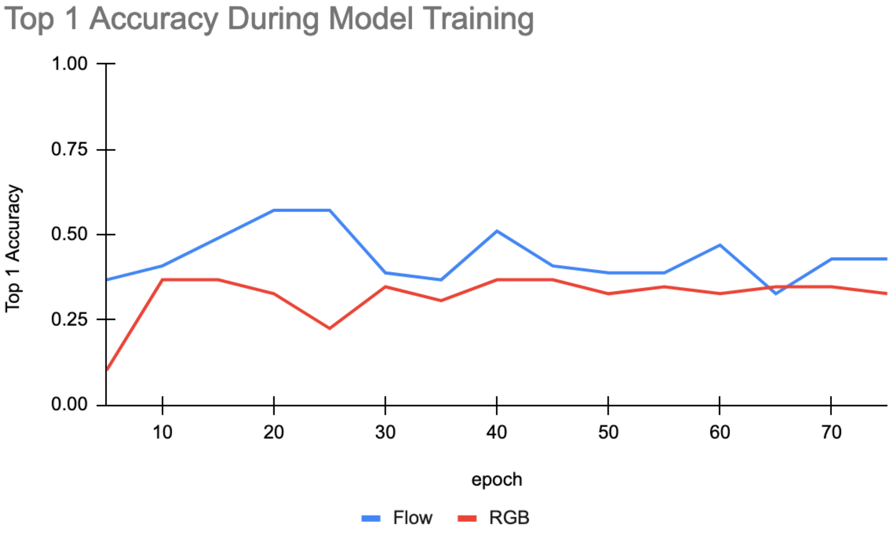

*Figure 24: Top 1 accuracy during training for RGB and Optical Flow networks.*

## Discussion & Future Work

These experiments outline the initial approach to applying advances in video behavior classification of human subjects to bumblebees. Specifically, this work established an initial dataset for deep learning models to train on bumblebee-specific behaviors. A total of 237 videos of five action classes were annotated for use in model training, validation, and testing. According to the preliminary results, the low accuracy of both network models in TSN indicate that further work is needed before automated behavior recognition can be used for bumblebee conservation efforts. Importantly, the temporal network component did achieve higher accuracy of 57.75% than the baseline accuracy of 35.2%. The spatial network on the other hand barely improved on the baseline accuracy with a top 1 accuracy of 38.77%. These experiments highlight the need for a more robust and balanced dataset that can facilitate the use of deep learning architectures.

As discussed earlier, automating the behavioral classification of bumblebee videos has numerous benefits for bumblebee conservation. Continuing to improve behavior classification for bumblebees is important for conservation efforts and can be approached from several directions. The first, and perhaps most obvious source of improvement is the bumblebee behavior dataset itself. As Beecology users continue to submit videos, the potential scale of an annotated dataset increases. While modern deep learning approaches require far less data than older approaches, models still require sufficient data to learn novel behaviors. At 237 video clips, the current dataset is small for a novel problem and would certainly benefit from additional annotated videos. Furthermore, augmenting the annotations with start and end times can produce a dataset useful for training temporal models that require less data. This would be useful both for utilizing more recently developed deep learning architectures as well as enabling benchmark comparisons with datasets like AVA and ActivityNet. Besides augmenting the bumblebee dataset, next steps could address the current imbalance in data for behavioral classes. Currently, nectar foraging makes up most of the existing videos. While resampling from such a small original dataset can certainly cause overfitting issues, rebalancing the dataset once more videos are added should improve the training of models.

Additional improvements to the bumblebee video behavior classification process can be made with additional resources. A dedicated VM or hardware setup could allow for the use of a docker image that avoids issues that occur with using a shared remote server. Not needing to identify different software configurations when software is updated for graphic cards are replaced would significantly speed up progress. Long term improvements could include training TSN from scratch on a bumblebee dataset rather than initializing weights from previous models. Using weights from previous models is generally a useful time-saving strategy as training a model from scratch is very computationally intensive. However, existing model weights are based on human datasets, and it is unknown how well models of human behavior transfer to bumblebee behavior.

# Chapter 6: References

1. Cameron, S. A. & Sadd, B. M. Global Trends in Bumble Bee Health. *Annual Review of Entomology* **65**, 209–232 (2020).
2. The Beecology Project. https://beecology.wpi.edu/website/home.
3. Cardinal, S. & Danforth, B. N. The Antiquity and Evolutionary History of Social Behavior in Bees. *PLoS One* **6**, e21086 (2011).
4. Honeybees vs bumblebees | Bumblebee Conservation Trust. https://web.archive.org/web/20150228004017/http://bumblebeeconservation.org/about-bees/faqs/honeybees-vs-bumblebees (2015).
5. Parrey, A. H. *et al.* Role of Bumblebees (Hymenoptera: Apidae) in Pollination of High Land Ecosystems: A Review. *Agricultural Reviews* (2021).
6. Gurel, F., Gosterit, A. & Eren, Ö. Life-cycle and foraging patterns of native Bombus terrestris (L.) (Hymenoptera, Apidae) in the Mediterranean region. *Insect. Soc.* **55**, 123–128 (2008).
7. Lifecycle - Bumblebee Conservation Trust. https://www.bumblebeeconservation.org/lifecycle/.
8. Dornhaus, A. & Chittka, L. Bumble bees (Bombus terrestris) store both food and information in honeypots. *Behav Ecol* **16**, 661–666 (2005).
9. Nayak, R. K., Rana, K., Bairwa, V. K. & Bharthi, V. D. A review on role of bumblebee pollination in fruits and vegetables. 7.
10. Steffan-Dewenter, I. & Westphal, C. The interplay of pollinator diversity, pollination services and landscape change. *Journal of Applied Ecology* **45**, 737–741 (2008).
11. Goulson, D., Lye, G. C. & Darvill, B. Decline and Conservation of Bumble Bees. *Annual Review of Entomology* **53**, 191–208 (2008).
12. Landi, P., Minoarivelo, H. O., Brännström, Å., Hui, C. & Dieckmann, U. Complexity and stability of ecological networks: a review of the theory. *Population Ecology* **60**, 319–345 (2018).
13. Memmott, J., Craze, P. G., Waser, N. M. & Price, M. V. Global warming and the disruption of plant–pollinator interactions. *Ecology Letters* **10**, 710–717 (2007).
14. Cameron, S. A. *et al.* Patterns of widespread decline in North American bumble bees. *PNAS* **108**, 662–667 (2011).
15. Goulson, D., Nicholls, E., Botías, C. & Rotheray, E. L. Bee declines driven by combined stress from parasites, pesticides, and lack of flowers. *Science* **347**, 1255957 (2015).
16. Biesmeijer, J. C. *et al.* Parallel Declines in Pollinators and Insect-Pollinated Plants in Britain and the Netherlands. *Science* **313**, 351–354 (2006).
17. Becher, M. A., Osborne, J. L., Thorbek, P., Kennedy, P. J. & Grimm, V. REVIEW: Towards a systems approach for understanding honeybee decline: a stocktaking and synthesis of existing models. *Journal of Applied Ecology* **50**, 868–880 (2013).
18. Henry, M. *et al.* Predictive systems models can help elucidate bee declines driven by multiple combined stressors. *Apidologie* **48**, 328–339 (2017).
19. Tapparo, A. *et al.* Assessment of the Environmental Exposure of Honeybees to Particulate Matter Containing Neonicotinoid Insecticides Coming from Corn Coated Seeds. *Environ. Sci. Technol.* **46**, 2592–2599 (2012).
20. Blacquière, T., Smagghe, G., van Gestel, C. A. M. & Mommaerts, V. Neonicotinoids in bees: a review on concentrations, side-effects and risk assessment. *Ecotoxicology* **21**, 973–992 (2012).

21. Lundin, O., Rundlöf, M., Smith, H. G., Fries, I. & Bommarco, R. Neonicotinoid Insecticides and Their Impacts on Bees: A Systematic Review of Research Approaches and Identification of Knowledge Gaps. *PLOS ONE* **10**, e0136928 (2015).
22. Wood, S. C., Kozii, I. V., Koziy, R. V., Epp, T. & Simko, E. Comparative chronic toxicity of three neonicotinoids on New Zealand packaged honey bees. *PLOS ONE* **13**, e0190517 (2018).
23. Cresswell, J. E. A demographic approach to evaluating the impact of stressors on bumble bee colonies. *Ecological Entomology* **42**, 221–229 (2017).
24. Czerwinski, M. & Sadd, B. Detrimental interactions of neonicotinoid pesticide exposure and bumblebee immunity. (2017).
25. Bryden, J., Gill, R. J., Mitton, R. A. A., Raine, N. E. & Jansen, V. A. A. Chronic sublethal stress causes bee colony failure. *Ecology Letters* **16**, 1463–1469.
26. Mobley, M. W. Examining the Potential Threat of Pesticide and Pathogen Exposure on Wild Bumble Bees: Proposed Lethal and Sublethal Mechanisms Contributing to Pollinator Decline. 274 (2017).
27. Piiroinen, S. & Goulson, D. Chronic neonicotinoid pesticide exposure and parasite stress differentially affects learning in honeybees and bumblebees. *Proceedings of the Royal Society B: Biological Sciences* **283**, 20160246 (2016).
28. van der Sluijs, J. P. *et al.* Neonicotinoids, bee disorders and the sustainability of pollinator services. *Current Opinion in Environmental Sustainability* **5**, 293–305 (2013).
29. Yang, E. C., Chuang, Y. C., Chen, Y. L. & Chang, L. H. Abnormal Foraging Behavior Induced by Sublethal Dosage of Imidacloprid in the Honey Bee (Hymenoptera: Apidae). *Journal of Economic Entomology* **101**, 1743–1748 (2008).
30. Han, P., Niu, C.-Y., Lei, C.-L., Cui, J.-J. & Desneux, N. Quantification of toxins in a Cry1Ac + CpTI cotton cultivar and its potential effects on the honey bee Apis mellifera L. *Ecotoxicology* **19**, 1452–1459 (2010).
31. Mommaerts, V. *et al.* Risk assessment for side-effects of neonicotinoids against bumblebees with and without impairing foraging behavior. *Ecotoxicology* **19**, 207 (2009).
32. Henry, M. *et al.* A Common Pesticide Decreases Foraging Success and Survival in Honey Bees. *Science* **336**, 348–350 (2012).
33. Wright, G. A., Softley, S. & Earnshaw, H. Low doses of neonicotinoid pesticides in food rewards impair short-term olfactory memory in foraging-age honeybees. *Scientific Reports* **5**, (2015).
34. Otterstatter, M. C., Gegear, R. J., Colla, S. R. & Thomson, J. D. Effects of parasitic mites and protozoa on the flower constancy and foraging rate of bumble bees. *Behav Ecol Sociobiol* **58**, 383–389 (2005).
35. Riddell, C. E. & Mallon, E. B. Insect psychoneuroimmunology: Immune response reduces learning in protein starved bumblebees (Bombus terrestris). *Brain, Behavior, and Immunity* **20**, 135–138 (2006).
36. Gegear, R. J., Otterstatter, M. C. & Thomson, J. D. Bumble-bee foragers infected by a gut parasite have an impaired ability to utilize floral information. *Proceedings of the Royal Society B: Biological Sciences* (2006) doi:10.1098/rspb.2005.3423.
37. Siviter, H., Koricheva, J., Brown, M. J. F. & Leadbeater, E. Quantifying the impact of pesticides on learning and memory in bees. *Journal of Applied Ecology* **55**, 2812–2821 (2018).

38. Stanley, D. A., Smith, K. E. & Raine, N. E. Bumblebee learning and memory is impaired by chronic exposure to a neonicotinoid pesticide. *Scientific Reports* **5**, (2015).
39. Mobley, M. W. & Gegear, R. J. Immune-cognitive system connectivity reduces bumblebee foraging success in complex multisensory floral environments. *Scientific Reports* **8**, 5953 (2018).
40. Gómez-Moracho, T., Heeb, P. & Lihoreau, M. Effects of parasites and pathogens on bee cognition. *Ecological Entomology* **42**, 51–64 (2017).
41. Decourtye, A. *et al.* Imidacloprid impairs memory and brain metabolism in the honeybee (Apis mellifera L.). *Pesticide Biochemistry and Physiology* **78**, 83–92 (2004).
42. Laycock, I., Lenthall, K. M., Barratt, A. T. & Cresswell, J. E. Effects of imidacloprid, a neonicotinoid pesticide, on reproduction in worker bumble bees (Bombus terrestris). *Ecotoxicology* **21**, 1937–1945 (2012).
43. Rumkee, J. C. O., Becher, M. A., Thorbek, P. & Osborne, J. L. Modeling Effects of Honeybee Behaviors on the Distribution of Pesticide in Nectar within a Hive and Resultant in-Hive Exposure. *Environ. Sci. Technol.* **51**, 6908–6917 (2017).
44. Schneider, C. W., Tautz, J., Grünewald, B. & Fuchs, S. RFID Tracking of Sublethal Effects of Two Neonicotinoid Insecticides on the Foraging Behavior of Apis mellifera. *PLOS ONE* **7**, e30023 (2012).
45. Tirado, W. R., Simon, G., Johnston, P. & Boer, P. Bees in decline Greenpeace Research Laboratories Technical Report (Review) 01/2013. 48 (2013).
46. Whitehorn, P. R., O'Connor, S., Wackers, F. L. & Goulson, D. Neonicotinoid Pesticide Reduces Bumble Bee Colony Growth and Queen Production. *Science* **336**, 351–352 (2012).
47. Willis Chan, D. S. & Raine, N. E. Population decline in a ground-nesting solitary squash bee (Eucera pruinosa) following exposure to a neonicotinoid insecticide treated crop (Cucurbita pepo). *Sci Rep* **11**, 4241 (2021).
48. Sargent, C. *et al.* Acute Imidacloprid Exposure Alters Mitochondrial Function in Bumblebee Flight Muscle and Brain. *Frontiers in Insect Science* **1**, (2021).
49. Tasman, K., Rands, S. A. & Hodge, J. J. L. The Neonicotinoid Insecticide Imidacloprid Disrupts Bumblebee Foraging Rhythms and Sleep. *iScience* **23**, (2020).
50. Goulson, D. Neonicotinoids impact bumblebee colony fitness in the field; a reanalysis of the UK's Food & Environment Research Agency 2012 experiment. *PeerJ* **3**, e854 (2015).
51. Li, Y., Miao, R. & Khanna, M. Neonicotinoids and decline in bird biodiversity in the United States. *Nat Sustain* **3**, 1027–1035 (2020).
52. Chittka, L. & Thomson, J. D. *Cognitive Ecology of Pollination: Animal Behaviour and Floral Evolution*. (Cambridge University Press, 2001).
53. Elzinga, J. A. *et al.* Time after time: flowering phenology and biotic interactions. *Trends in Ecology & Evolution* **22**, 432–439 (2007).
54. Chittka, L., Thomson, J. D. & Waser, N. M. Flower Constancy, Insect Psychology, and Plant Evolution. *Naturwissenschaften* **86**, 361–377 (1999).
55. Bronstein, J. L., Alarcón, R. & Geber, M. The evolution of plant–insect mutualisms. *New Phytologist* **172**, 412–428 (2006).
56. Bond, W. J., Lawton, J. H. & May, R. M. Do mutualisms matter? Assessing the impact of pollinator and disperser disruption on plant extinction. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* **344**, 83–90 (1994).
57. Berec, L. Mixed encounters, limited perception and optimal foraging. *Bull. Math. Biol.* **62**, 849–868 (2000).

58. Dunlap, A. S., Papaj, D. R. & Dornhaus, A. Sampling and tracking a changing environment: persistence and reward in the foraging decisions of bumblebees. *Interface Focus* **7**, 20160149 (2017).
59. Goulson, D. Foraging strategies of insects for gathering nectar and pollen, and implications for plant ecology and evolution. *Perspectives in Plant Ecology, Evolution and Systematics* **2**, 185–209 (1999).
60. Raine, N. E. & Chittka, L. The Adaptive Significance of Sensory Bias in a Foraging Context: Floral Colour Preferences in the Bumblebee Bombus terrestris. *PLOS ONE* **2**, e556 (2007).
61. Chittka, L. & Raine, N. E. Recognition of flowers by pollinators. *Current Opinion in Plant Biology* **9**, 428–435 (2006).
62. Bartomeus, I. *et al.* Climate-associated phenological advances in bee pollinators and bee-pollinated plants. *PNAS* **108**, 20645–20649 (2011).
63. Dall, S. R. X., Giraldeau, L.-A., Olsson, O., McNamara, J. M. & Stephens, D. W. Information and its use by animals in evolutionary ecology. *Trends in Ecology & Evolution* **20**, 187–193 (2005).
64. Piao, S. *et al.* Plant phenology and global climate change: Current progresses and challenges. *Global Change Biology* **25**, 1922–1940 (2019).
65. Hendriksma, H. P., Toth, A. L. & Shafir, S. Individual and Colony Level Foraging Decisions of Bumble Bees and Honey Bees in Relation to Balancing of Nutrient Needs. *Front. Ecol. Evol.* **7**, (2019).
66. Raine, N. E. & Chittka, L. The correlation of learning speed and natural foraging success in bumble-bees. *Proc Biol Sci* **275**, 803–808 (2008).
67. Waser, N. M. Flower Constancy: Definition, Cause, and Measurement. *The American Naturalist* **127**, (1986).
68. Rademaker, M. C. J., De Jong, T. J. & Klinkhamer, P. G. L. Pollen Dynamics of Bumble-Bee Visitation on Echium vulgare. *Functional Ecology* **11**, 554–563 (1997).
69. Midgley, J. J. *et al.* How important is biotic pollination and dispersal to the success of the angiosperms? *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* **333**, 209–215 (1991).
70. Baracchi, D. Cognitive ecology of pollinators and the main determinants of foraging plasticity. *Current Zoology* **65**, 421–424 (2019).
71. Spaethe, J., Tautz, J. & Chittka, L. Visual constraints in foraging bumblebees: Flower size and color affect search time and flight behavior. *PNAS* **98**, 3898–3903 (2001).
72. Goulson, D. *et al.* Choosing rewarding flowers; perceptual limitations and innate preferences influence decision making in bumblebees and honeybees. *Behav Ecol Sociobiol* **61**, 1523–1529 (2007).
73. Keasar, T., Motro, U. & Shmida, A. Temporal reward variability promotes sampling of a new flower type by bumblebees. *Animal Behaviour* **86**, 747–753 (2013).
74. Shettleworth, S. J. Cognition, Evolution, and Behavior. 2 (1998).
75. Giurfa, M. Cognition with few neurons: higher-order learning in insects. *Trends in Neurosciences* **36**, 285–294 (2013).
76. Webb, B. Cognition in insects. *Phil. Trans. R. Soc. B* **367**, 2715–2722 (2012).
77. Menzel, R. The honeybee as a model for understanding the basis of cognition. *Nat Rev Neurosci* **13**, 758–768 (2012).
78. Menzel, R. & Giurfa, M. Dimensions of Cognition in an Insect, the Honeybee. *Behavioral and Cognitive Neuroscience Reviews* (2006).

79. Heinrich, B. Resource Heterogeneity and Patterns of Movement in Foraging Bumblebees. *Oecologia* **40**, 235–245 (1979).
80. Allen, T., Cameron, S., McGinley, R. & Heinrich, B. The Role of Workers and New Queens in the Ergonomics of a Bumblebee Colony (Hymenoptera: Apoidea). *Journal of the Kansas Entomological Society* **51**, 329–342 (1978).
81. Willmer, P. G. & Stone, G. N. *Advances in the Study of Behavior*. (Elsevier, 2005).
82. Raine, N. E. & Chittka, L. Flower Constancy and Memory Dynamics in Bumblebees (Hymenoptera: Apidae: Bombus). 21 (2006).
83. Menzel, R. Searching for the Memory Trace in a Mini-Brain, the Honeybee. *Learning & Memory* **8**, 53–62 (2001).
84. Briscoe, A. D. & Chittka, L. The Evolution of Color Vision in Insects. *Annual Review of Entomology* **46**, 471–510 (2001).
85. Scheiner, R., Page, R. E. & Erber, J. The Effects of Genotype, Foraging Role, and Sucrose Responsiveness on the Tactile Learning Performance of Honey Bees (Apis mellifera L.). *Neurobiology of Learning and Memory* **76**, 138–150 (2001).
86. Ings, T. C. & Chittka, L. Predator crypsis enhances behaviourally mediated indirect effects on plants by altering bumblebee foraging preferences. *Proc Biol Sci* **276**, 2031–2036 (2009).
87. Ings, T. C., Raine, N. E. & Chittka, L. A population comparison of the strength and persistence of innate colour preference and learning speed in the bumblebee Bombus terrestris. *2009*.
88. Muth, F., Cooper, T. R., Bonilla, R. F. & Leonard, A. S. A novel protocol for studying bee cognition in the wild. *Methods in Ecology and Evolution* **9**, 78–87 (2017).
89. Morand-Ferron, J., Hamblin, S., Cole, E. F., Aplin, L. M. & Quinn, J. L. Taking the Operant Paradigm into the Field: Associative Learning in Wild Great Tits. *PLOS ONE* **10**, e0133821 (2015).
90. Morand-Ferron, J., Cole, E. F. & Quinn, J. L. Studying the evolutionary ecology of cognition in the wild: a review of practical and conceptual challenges. *Biological Reviews* **91**, 367–389 (2016).
91. Rowe, C. & Healy, S. D. Measuring variation in cognition. *Behav Ecol* **25**, 1287–1292 (2014).
92. Giurfa, M. Behavioral and neural analysis of associative learning in the honeybee: a taste from the magic well. *J Comp Physiol A* **193**, 801–824 (2007).
93. Baddeley, A. D. Working Memory. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* **302**, 311–324 (1983).
94. Matzel, L. D. & Kolata, S. Selective attention, working memory, and animal intelligence. *Neuroscience & Biobehavioral Reviews* **34**, 23–30 (2010).
95. Zhang, S., Bock, F., Si, A., Tautz, J. & Srinivasan, M. V. Visual working memory in decision making by honey bees. *PNAS* **102**, 5250–5255 (2005).
96. Vandierendonck, A. Role of Working Memory in Task Switching. *PSYCHOL BELG* **52**, 229 (2012).
97. Bitterman, M. E., Menzel, R., Fietz, A. & Schäfer, S. Classical conditioning of proboscis extension in honeybees (Apis mellifera). *Journal of Comparative Psychology* **97**, 107–119 (1983).
98. Takeda, K. Classical conditioned response in the honey bee. *Journal of Insect Physiology* **6**, 168–179 (1961).
99. Menzel, R. Memory dynamics in the honeybee. *J Comp Physiol A* **185**, 323–340 (1999).

100. Becher, M. A. *et al.* Bumble-BEEHAVE: A systems model for exploring multifactorial causes of bumblebee decline at individual, colony, population and community level. *Journal of Applied Ecology* **0**, (2018).
101. DeAngelis, D. L. & Grimm, V. Individual-based models in ecology after four decades. *F1000Prime Rep* **6**, (2014).
102. Mortensen, L. O., Chudzinska, M. E., Slabbekoorn, H. & Thomsen, F. Agent-based models to investigate sound impact on marine animals: bridging the gap between effects on individual behaviour and population level consequences. *Oikos* **130**, 1074–1086 (2021).
103. Retzlaff, C. O., Ziefle, M. & Calero Valdez, A. The History of Agent-Based Modeling in the Social Sciences. in *Digital Human Modeling and Applications in Health, Safety, Ergonomics and Risk Management. Human Body, Motion and Behavior* (ed. Duffy, V. G.) 304–319 (Springer International Publishing, 2021). doi:10.1007/978-3-030-77817-0_22.
104. Tesfatsion, L. Agent-Based Computational Economics: Growing Economies From the Bottom Up. *Artificial Life* **8**, 55–82 (2002).
105. McLane, A. J., Semeniuk, C., McDermid, G. J. & Marceau, D. J. The role of agent-based models in wildlife ecology and management. *Ecological Modelling* **222**, 1544–1556 (2011).
106. Stillman, R. A., Railsback, S. F., Giske, J., Berger, U. & Grimm, V. Making Predictions in a Changing World: The Benefits of Individual-Based Ecology. *BioScience* **65**, 140–150 (2015).
107. Grimm, V. & Berger, U. Structural realism, emergence, and predictions in next-generation ecological modelling: Synthesis from a special issue. *Ecological Modelling* **326**, 177–187 (2016).
108. Grimm, V. & Railsback, S. F. Pattern-oriented modelling: a 'multi-scope' for predictive systems ecology. *Phil. Trans. R. Soc. B* **367**, 298–310 (2012).
109. Qu, H., Seifan, T., Tielbörger, K. & Seifan, M. A spatially explicit agent-based simulation platform for investigating effects of shared pollination service on ecological communities. *Simulation Modelling Practice and Theory* **37**, 107–124 (2013).
110. Olsson, O., Bolin, A., Smith, H. G. & Lonsdorf, E. V. Modeling pollinating bee visitation rates in heterogeneous landscapes from foraging theory. *Ecological Modelling* **316**, 133–143 (2015).
111. Crone, E. E. & Williams, N. M. Bumble bee colony dynamics: quantifying the importance of land use and floral resources for colony growth and queen production. *Ecology Letters* **19**, 460–468 (2016).
112. Dornhaus, A. & Chittka, L. Food alert in bumblebees (Bombus terrestris): possible mechanisms and evolutionary implications. (2001).
113. Dornhaus, A. & Chittka, L. Information flow and regulation of foraging activityin bumble bees (Bombus spp.). *Apidologie* **35**, 183–192 (2004).
114. Saleh, N., Ohashi, K., Thomson, J. D. & Chittka, L. Facultative use of the repellent scent mark in foraging bumblebees: complex versus simple flowers. *Animal Behaviour* **71**, 847–854 (2006).
115. Molet, M., Chittka, L. & Raine, N. E. How floral odours are learned inside the bumblebee (Bombus terrestris) nest. *Naturwissenschaften* **96**, 213–219 (2009).
116. Twiston-Davies, G., Becher, M. A. & Osborne, J. L. BEE-STEWARD: A research and decision-support software for effective land management to promote bumblebee populations. *Methods in Ecology and Evolution* **12**, 1809–1815 (2021).

117. Samuelson, E. E. W., Chen-Wishart, Z. P., Gill, R. J. & Leadbeater, E. Effect of acute pesticide exposure on bee spatial working memory using an analogue of the radial-arm maze. *Scientific Reports* **6**, srep38957 (2016).
118. Versypt, A., Crall, J. & Dey, B. BeeNestABM: An open-source agent-based model of spatiotemporal distribution of bumblebees in nests. *Journal of Open Source Software* **3**, 718 (2018).
119. Bryden, J., Gill, R. J., Mitton, R. A. A., Raine, N. E. & Jansen, V. A. A. Chronic sublethal stress causes bee colony failure. *Ecology Letters* **16**, 1463–1469 (2013).
120. Banks, H. T. *et al.* Modeling bumble bee population dynamics with delay differential equations. *Ecological Modelling* **351**, 14–23 (2017).
121. Banks, J. E., Banks, H. T., Myers, N., Laubmeier, A. N. & Bommarco, R. Lethal and sublethal effects of toxicants on bumble bee populations: a modelling approach. *Ecotoxicology* **29**, 237–245 (2020).
122. Biernaskie, J. M., Walker, S. C. & Gegear, R. J. Bumblebees Learn to Forage like Bayesians. *The American Naturalist* **174**, 413–423 (2009).
123. Häussler, J., Sahlin, U., Baey, C., Smith, H. G. & Clough, Y. Pollinator population size and pollination ecosystem service responses to enhancing floral and nesting resources. *Ecology and Evolution* **7**, 1898–1908 (2016).
124. McKinley, D. C. *et al.* Citizen science can improve conservation science, natural resource management, and environmental protection. *Biological Conservation* **208**, 15–28 (2017).
125. Ballard, H. L. *et al.* Contributions to conservation outcomes by natural history museum-led citizen science: Examining evidence and next steps. *Biological Conservation* **208**, 87–97 (2017).
126. Bonney, R. *et al.* Next Steps for Citizen Science. *Science* **343**, 1436–1437 (2014).
127. Conrad, C. C. & Hilchey, K. G. A review of citizen science and community-based environmental monitoring: issues and opportunities. *Environmental Monitoring and Assessment* **176**, 273–291 (2011).
128. Falk, S. *et al.* Evaluating the ability of citizen scientists to identify bumblebee (Bombus) species. *PLoS ONE* **14**, e0218614 (2019).
129. Grimm, V. *et al.* The ODD Protocol for Describing Agent-Based and Other Simulation Models: A Second Update to Improve Clarity, Replication, and Structural Realism. *JASSS* **23**, 7 (2020).
130. Wilensky, U. NetLogo. http://ccl.northwestern.edu/netlogo/ (1999).
131. Vaudo, A. D., Farrell, L. M., Patch, H. M., Grozinger, C. M. & Tooker, J. F. Consistent pollen nutritional intake drives bumble bee (Bombus impatiens) colony growth and reproduction across different habitats. *Ecology and Evolution* **8**, 5765–5776 (2018).
132. Heinrich, B. The Foraging Specializations of Individual Bumblebees. *Ecological Monographs* **46**, 105–128 (1976).
133. Castellanos, M. C., Wilson, P. & Thomson, J. D. Dynamic nectar replenishment in flowers of Penstemon (Scrophulariaceae). *American Journal of Botany* **89**, 111–118 (2002).
134. Metz, J. *et al.* Plant survival in relation to seed size along environmental gradients: a long-term study from semi-arid and Mediterranean annual plant communities. *Journal of Ecology* **98**, 697–704 (2010).
135. Cane, J. H. & Dunne, R. Generalist Bees Pollinate Red-flowered Penstemon eatonii: Duality in the Hummingbird Pollination Syndrome. *amid* **171**, 365–370 (2014).

136. Castellanos, M. C., Wilson, P. & Thomson, J. D. Pollen Transfer by Hummingbirds and Bumblebees, and the Divergence of Pollination Modes in Penstemon. *Evolution* **57**, 2742–2752 (2003).

137. Mondal, A. K. & Mandal, S. Pollen production in some terrestrial angiosperms. *Current Science* **74**, 906–910 (1998).

138. Dunlap, A. S. & Stephens, D. W. Tracking a changing environment: optimal sampling, adaptive memory and overnight effects. *Behavioural Processes* **89**, 86–94 (2012).

139. Charnov, E. L. Optimal foraging, the marginal value theorem. *Theoretical Population Biology* **9**, 129–136 (1976).

140. Pyke, G. H. Optimal foraging theory: a critical review. *Annual review of ecology and systematics* **15**, 523–575 (1984).

141. Grimm, V. *et al.* A standard protocol for describing individual-based and agent-based models. *Ecological Modelling* **198**, 115–126 (2006).

142. Grimm, V. *et al.* The ODD protocol: A review and first update. *Ecological Modelling* **221**, 2760–2768 (2010).

143. Folly, A. J., Koch, H., Stevenson, P. C. & Brown, M. J. F. Larvae act as a transient transmission hub for the prevalent bumblebee parasite Crithidia bombi. *Journal of Invertebrate Pathology* **148**, 81–85 (2017).

144. Colla, S. R., Otterstatter, M. C., Gegear, R. J. & Thomson, J. D. Plight of the bumble bee: Pathogen spillover from commercial to wild populations. *Biological Conservation* **129**, 461–467 (2006).

145. Gillespie, S. Factors affecting parasite prevalence among wild bumblebees. *Ecological Entomology* **35**, 737–747 (2010).

146. Goulson, D. REVIEW: An overview of the environmental risks posed by neonicotinoid insecticides. *Journal of Applied Ecology* **50**, 977–987 (2013).

147. Krupke, C. H., Hunt, G. J., Eitzer, B. D., Andino, G. & Given, K. Multiple Routes of Pesticide Exposure for Honey Bees Living Near Agricultural Fields. *PLOS ONE* **7**, e29268 (2012).

148. Dolezal, A. G. *et al.* Honey Bee Viruses in Wild Bees: Viral Prevalence, Loads, and Experimental Inoculation. *PLOS ONE* **11**, e0166190 (2016).

149. Hladik, M. L., Main, A. R. & Goulson, D. Environmental Risks and Challenges Associated with Neonicotinoid Insecticides. *Environ. Sci. Technol.* **52**, 3329–3335 (2018).

150. Banaszak-Cibicka, W. & Żmihorski, M. Wild bees along an urban gradient: winners and losers. *J Insect Conserv* **16**, 331–343 (2012).

151. Traveset, A., Tur, C. & Eguíluz, V. M. Plant survival and keystone pollinator species in stochastic coextinction models: role of intrinsic dependence on animal-pollination. *Sci Rep* **7**, 6915 (2017).

152. Vieira, M. C. & Almeida-Neto, M. A simple stochastic model for complex coextinctions in mutualistic networks: robustness decreases with connectance. *Ecology Letters* **18**, 144–152 (2015).

153. Newton, I. The recent declines of farmland bird populations in Britain: an appraisal of causal factors and conservation actions. *Ibis* **146**, 579–600 (2004).

154. Titulaer, M. *et al.* Molecular analysis of stomach contents reveals important grass seeds in the winter diet of Baird's and Grasshopper sparrows, two declining grassland bird species. *PLOS ONE* **12**, e0189695 (2017).

155. Tur, C., Vigalondo, B., Trøjelsgaard, K., Olesen, J. M. & Traveset, A. Downscaling pollen-transport networks to the level of individuals. *J Anim Ecol* **83**, 306–317 (2014).
156. Gegear, R. J. & Laverty, T. M. Flower constancy in bumblebees: a test of the trait variability hypothesis. *Animal Behaviour* **69**, 939–949 (2005).
157. Ellner, S. P., Ng, W. H. & Myers, C. R. Individual Specialization and Multihost Epidemics: Disease Spread in Plant-Pollinator Networks. *The American Naturalist* E000–E000 (2020) doi:10.1086/708272.
158. Morán-López, T. *et al.* Can network metrics predict vulnerability and species roles in bird-dispersed plant communities? Not without behaviour. *Ecol Lett* **23**, 348–358 (2020).
159. Strang, C. G. & Sherry, D. F. Serial reversal learning in bumblebees (Bombus impatiens). *Anim Cogn* **17**, 723–734 (2014).
160. Maharaj, G., Horack, P., Yoder, M. & Dunlap, A. S. Influence of preexisting preference for color on sampling and tracking behavior in bumble bees. *Behav Ecol* **30**, 150–158 (2019).
161. Chittka, L. & Niven, J. Are Bigger Brains Better? *Current Biology* **19**, R995–R1008 (2009).
162. Chittka, L. Bee cognition. *Current Biology* **27**, R1049–R1053 (2017).
163. Greggers, U. & Menzel, R. Memory Dynamics and Foraging Strategies of Honeybees. *Behavioral Ecology and Sociobiology* **32**, 17–29 (1993).
164. Sih, A., Trimmer, P. C. & Ehlman, S. M. A conceptual framework for understanding behavioral responses to HIREC. *Current Opinion in Behavioral Sciences* **12**, 109–114 (2016).
165. Perry, C. J., Barron, A. B. & Chittka, L. The frontiers of insect cognition. *Current Opinion in Behavioral Sciences* **16**, 111–118 (2017).
166. Ragozzino, M. E. The Contribution of the Medial Prefrontal Cortex, Orbitofrontal Cortex, and Dorsomedial Striatum to Behavioral Flexibility. *Annals of the New York Academy of Sciences* **1121**, 355–375 (2007).
167. Coppens, C. M., de Boer, S. F. & Koolhaas, J. M. Coping styles and behavioural flexibility: towards underlying mechanisms. *Philosophical Transactions of the Royal Society B: Biological Sciences* **365**, 4021–4028 (2010).
168. Lea, S. E. G., Chow, P. K. Y., Leaver, L. A. & McLaren, I. P. L. Behavioral flexibility: A review, a model, and some exploratory tests. *Learn Behav* **48**, 173–187 (2020).
169. Strang, C. Behavioural Flexibility in Bumblebees (Bombus impatiens). 165.
170. Audet, J.-N. & Lefebvre, L. What's flexible in behavioral flexibility? *Behavioral Ecology* **28**, 943–947 (2017).
171. Sih, A. Understanding variation in behavioural responses to human-induced rapid environmental change: a conceptual overview. *Animal Behaviour* **85**, 1077–1088 (2013).
172. Sih, A., Ferrari, M. C. O. & Harris, D. J. Evolution and behavioural responses to human-induced rapid environmental change. *Evolutionary Applications* **4**, 367–387 (2011).
173. Hethcoat, M. G. & Chalfoun, A. D. Towards a mechanistic understanding of human-induced rapid environmental change: a case study linking energy development, nest predation and predators. *Journal of Applied Ecology* **52**, 1492–1499 (2015).
174. Gill, K. S. *et al.* Experimental Test-Bed for Bumblebee-Inspired Channel Selection in an Ad-Hoc Network. in *2018 IEEE 88th Vehicular Technology Conference (VTC-Fall)* 1–5 (2018). doi:10.1109/VTCFall.2018.8690978.

175. Hwang, J., Bose, N., Nguyen, H. D. & Williams, G. Acoustic Search and Detection of Oil Plumes Using an Autonomous Underwater Vehicle. *Journal of Marine Science and Engineering* **8**, 618 (2020).
176. Tello-Ramos, M. C., Branch, C. L., Kozlovsky, D. Y., Pitera, A. M. & Pravosudov, V. V. Spatial memory and cognitive flexibility trade-offs: to be or not to be flexible, that is the question. *Animal Behaviour* **147**, 129–136 (2019).
177. Sherry, D. F. & Strang, C. G. Contrasting styles in cognition and behaviour in bumblebees and honeybees. *Behavioural Processes* **117**, 59–69 (2015).
178. Bond, A. B., Kamil, A. C. & Balda, R. P. Serial reversal learning and the evolution of behavioral flexibility in three species of North American corvids (Gymnorhinus cyanocephalus, Nucifraga columbiana, Aphelocoma californica). *Journal of Comparative Psychology* **121**, 372–379 (2007).
179. Giurfa, M. Learning and cognition in insects. *WIREs Cognitive Science* **6**, 383–395 (2015).
180. Mackintosh, N. J. & Cauty, A. Spatial reversal learning in rats, pigeons, and goldfish. *Psychon Sci* **22**, 281–282 (1971).
181. Konstantinidis, E., Harman, J. L. & Gonzalez, C. Patterns of choice adaptation in dynamic risky environments. Preprint at https://doi.org/10.31234/osf.io/f34qb (2020).
182. Lejarraga, T., Lejarraga, J. & Gonzalez, C. Decisions from experience: How groups and individuals adapt to change. *Mem Cogn* **42**, 1384–1397 (2014).
183. Hertwig, R., Barron, G., Weber, E. U. & Erev, I. Decisions from Experience and the Effect of Rare Events in Risky Choice. *Psychol Sci* **15**, 534–539 (2004).
184. McCormick, E. N., Cheyette, S. J. & Gonzalez, C. Choice adaptation to changing environments: trends, feedback, and observability of change. *Mem Cogn* (2022) doi:10.3758/s13421-022-01313-2.
185. Brunswik, E. Probability as a determiner of rat behavior. *Journal of Experimental Psychology* **25**, 175–197 (1939).
186. Lloyd, K. & Leslie, D. S. Context-dependent decision-making: a simple Bayesian model. *Journal of The Royal Society Interface* **10**, 20130069 (2013).
187. Pulliam, H. R. On the Theory of Optimal Diets. *The American Naturalist* **108**, 59–74 (1974).
188. Zhang, F. & Hui, C. Recent experience-driven behaviour optimizes foraging. *Animal Behaviour* **88**, 13–19 (2014).
189. McNamara, J. & Houston, A. The application of statistical decision theory to animal behaviour. *Journal of Theoretical Biology* **85**, 673–690 (1980).
190. Dyer, A. G., Dorin, A., Reinhardt, V., Garcia, J. E. & Rosa, M. G. P. Bee reverse-learning behavior and intra-colony differences: Simulations based on behavioral experiments reveal benefits of diversity. *Ecological Modelling* **277**, 119–131 (2014).
191. MaBouDi, H., Solvi, C. & Chittka, L. Bumblebees Learn a Relational Rule but Switch to a Win-Stay/Lose-Switch Heuristic After Extensive Training. *Front Behav Neurosci* **14**, 137 (2020).
192. Bélisle, C. & Cresswell, J. The effects of a limited memory capacity on foraging behavior. *Theoretical population biology* **52**, 78–90 (1997).
193. Keasar, T., Rashkovich, E., Cohen, D. & Shmida, A. Bees in two-armed bandit situations: foraging choices and possible decision mechanisms. *Behavioral Ecology* **13**, 757–765 (2002).

194. Pyke, G. H. Optimal foraging: Movement patterns of bumblebees between inflorescences. *Theoretical Population Biology* **13**, 72–98 (1978).
195. Gigerenzer, G. & Todd, P. M. *Simple heuristics that make us smart*. (Oxford University Press, 1999).
196. Bílek, J., Nedoma, J. & Jirásek, M. REPRESENTATIVENESS HEURISTICS: A LITERATURE REVIEW OF ITS IMPACTS ON THE QUALITY OF DECISION-MAKING. 10.
197. Buckmann, M. & Şimşek, Ö. Decision Heuristics for Comparison:How Good Are They? in *Proceedings of the NIPS 2016 Workshop on Imperfect Decision Makers* 1–11 (PMLR, 2017).
198. Why Heuristics Work - Gerd Gigerenzer, 2008. https://journals.sagepub.com/doi/10.1111/j.1745-6916.2008.00058.x.
199. *Heuristics: The Foundations of Adaptive Behavior*. (Oxford University Press, 2011). doi:10.1093/acprof:oso/9780199744282.001.0001.
200. Hutchinson, J. M. C. & Gigerenzer, G. Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behavioural Processes* **69**, 97–124 (2005).
201. Hau, R., Pleskac, T. J., Kiefer, J. & Hertwig, R. The description–experience gap in risky choice: the role of sample size and experienced probabilities. *Journal of Behavioral Decision Making* **21**, 493–518 (2008).
202. Averell, L. & Heathcote, A. The form of the forgetting curve and the fate of memories. *Journal of Mathematical Psychology* **55**, 25–35 (2011).
203. Wixted, J. T. & Ebbesen, E. B. Genuine power curves in forgetting: A quantitative analysis of individual subject forgetting functions. *Memory & Cognition* **25**, 731–739 (1997).
204. Rescorla, R. & Wagner, A. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement (1972) Classical Conditioning II: Current Research and Theory. *Current Research and Theory* 64–99 (1998).
205. Niv, Y., Joel, D., Meilijson, I. & Ruppin, E. Evolution of Reinforcement Learning in Uncertain Environments: A Simple Explanation for Complex Foraging Behaviors. 20.
206. Frankenhuis, W. E., Panchanathan, K. & Barto, A. G. Enriching behavioral ecology with reinforcement learning methods. *Behavioural Processes* **161**, 94–100 (2019).
207. Yechiam, E. & Busemeyer, J. R. Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review* **12**, 387–402 (2005).
208. MaBouDi, H., Solvi, C. & Chittka, L. Bumblebees Learn a Relational Rule but Switch to a Win-Stay/Lose-Switch Heuristic After Extensive Training. *Front. Behav. Neurosci.* **14**, (2020).
209. Chittka, L. Sensorimotor learning in bumblebees: long-term retention and reversal training. *undefined* (1998).
210. Worden, B. D., Skemp, A. K. & Papaj, D. R. Learning in two contexts: the effects of interference and body size in bumblebees. *Journal of Experimental Biology* **208**, 2045–2053 (2005).
211. Stonedahl, F. & Wilensky, U. Finding Forms of Flocking: Evolutionary Search in ABM Parameter-Spaces. in *Multi-Agent-Based Simulation XI* (eds. Bosse, T., Geller, A. & Jonker, C. M.) 61–75 (Springer, 2011). doi:10.1007/978-3-642-18345-4_5.

212. rakow, T. & Miler, K. Doomed to repeat the successes of the past: History is best forgotten for repeated choices with nonstationary payoffs. *Memory & Cognition* **37**, 985–1000 (2009).
213. Rodríguez-Moreno, I., Martínez-Otzeta, J. M., Sierra, B., Rodriguez, I. & Jauregi, E. Video Activity Recognition: State-of-the-Art. *Sensors* **19**, 3160 (2019).
214. Guo, Y. *et al.* Deep learning for visual understanding: A review. *Neurocomputing* **187**, 27–48 (2016).
215. Herath, S., Harandi, M. & Porikli, F. Going deeper into action recognition: A survey. *Image and Vision Computing* **60**, 4–21 (2017).
216. Guo, K., Ishwar, P. & Konrad, J. Action Recognition Using Sparse Representation on Covariance Manifolds of Optical Flow. in *2010 7th IEEE International Conference on Advanced Video and Signal Based Surveillance* 188–195 (2010). doi:10.1109/AVSS.2010.71.
217. Kumar, S. S. & John, M. Human activity recognition using optical flow based feature set. in *2016 IEEE International Carnahan Conference on Security Technology (ICCST)* 1–5 (2016). doi:10.1109/CCST.2016.7815694.
218. Feng Niu & Abdel-Mottaleb, M. HMM-Based Segmentation and Recognition of Human Activities from Video Sequences. in *2005 IEEE International Conference on Multimedia and Expo* 804–807 (2005). doi:10.1109/ICME.2005.1521545.
219. Raman, N. & Maybank, S. J. Activity recognition using a supervised non-parametric hierarchical HMM. *Neurocomputing* **199**, 163–177 (2016).
220. Núñez, J. C., Cabido, R., Pantrigo, J. J., Montemayor, A. S. & Vélez, J. F. Convolutional Neural Networks and Long Short-Term Memory for skeleton-based human activity and hand gesture recognition. *Pattern Recognition* **76**, 80–94 (2018).
221. Sadanand, S. & Corso, J. J. Action bank: A high-level representation of activity in video. in *2012 IEEE Conference on Computer Vision and Pattern Recognition* 1234–1241 (2012). doi:10.1109/CVPR.2012.6247806.
222. Ng, J. Y.-H. & Davis, L. S. Temporal Difference Networks for Video Action Recognition. in *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)* 1587–1596 (2018). doi:10.1109/WACV.2018.00176.
223. Zhu, Y. *et al.* A Comprehensive Study of Deep Video Action Recognition. *arXiv:2012.06567 [cs]* (2020).
224. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T. & Serre, T. HMDB: A large video database for human motion recognition. in *2011 International Conference on Computer Vision* 2556–2563 (2011). doi:10.1109/ICCV.2011.6126543.
225. Soomro, K., Zamir, A. R. & Shah, M. UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. (2012) doi:10.48550/ARXIV.1212.0402.
226. Heilbron, F. C., Escorcia, V., Ghanem, B. & Niebles, J. C. ActivityNet: A large-scale video benchmark for human activity understanding. in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 961–970 (IEEE, 2015). doi:10.1109/CVPR.2015.7298698.
227. Gu, C. *et al.* AVA: A Video Dataset of Spatio-Temporally Localized Atomic Visual Actions. in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* 6047–6056 (2018). doi:10.1109/CVPR.2018.00633.

228. Monfort, M. *et al.* Moments in Time Dataset: one million videos for event understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1–1 (2019) doi:10.1109/TPAMI.2019.2901464.
229. Abu-El-Haija, S. *et al.* YouTube-8M: A Large-Scale Video Classification Benchmark. Preprint at http://arxiv.org/abs/1609.08675 (2016).
230. Simonyan, K. & Zisserman, A. Two-Stream Convolutional Networks for Action Recognition in Videos. in *NIPS* (2014).
231. Zhang, K., Li, D., Huang, J. & Chen, Y. Automated Video Behavior Recognition of Pigs Using Two-Stream Convolutional Networks. *Sensors* **20**, 1085 (2020).
232. Carreira, J. & Zisserman, A. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 4724–4733 (2017). doi:10.1109/CVPR.2017.502.
233. Lin, J., Gan, C. & Han, S. TSM: Temporal Shift Module for Efficient Video Understanding. in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)* 7082–7092 (2019). doi:10.1109/ICCV.2019.00718.
234. Li, Z., Gavrilyuk, K., Gavves, E., Jain, M. & Snoek, C. G. M. VideoLSTM convolves, attends and flows for action recognition. *Computer Vision and Image Understanding* **166**, 41–50 (2018).
235. Contributors, Mma. OpenMMLab's Next Generation Video Understanding Toolbox and Benchmark. https://github.com/open-mmlab/mmaction2 (2020).
236. Li, W., Swetha, S. & Shah, D. M. Wildlife Action Recognition using Deep Learning. 3.
237. Schütz, A. K. *et al.* Computer Vision for Detection of Body Posture and Behavior of Red Foxes. *Animals* **12**, 233 (2022).
238. Sclocco, A., Ong, S. J. Y., Aung, S. Y. P. & Teseo, S. *Integrating real-time data analysis into automatic tracking of social insect behavior*. http://biorxiv.org/lookup/doi/10.1101/2020.11.03.366195 (2020) doi:10.1101/2020.11.03.366195.
239. Rossetti, B. J., Dynes, T., Brosi, B., Roode, J. C. de & Kong, J. GRAPHITE: A graphical environment for scalable in situ video tracking of moving insects. *Methods in Ecology and Evolution* **9**, 956–964 (2018).
240. Woodgate, J. L., Makinson, J. C., Lim, K. S., Reynolds, A. M. & Chittka, L. Life-Long Radar Tracking of Bumblebees. *PLOS ONE* **11**, e0160333 (2016).
241. Michener, C. D. *The Bees of the World*. (JHU Press, 2000).
242. Goulson, D. Foraging strategies of insects for gathering nectar and pollen, and implications for plant ecology and evolution. *Perspectives in Plant Ecology, Evolution and Systematics* **2**, 185–209 (1999).
243. Konzmann, S. & Lunau, K. Divergent Rules for Pollen and Nectar Foraging Bumblebees – A Laboratory Study with Artificial Flowers Offering Diluted Nectar Substitute and Pollen Surrogate. *PLOS ONE* **9**, e91900 (2014).
244. Goulson, D. *et al.* Effects of land use at a landscape scale on bumblebee nest density and survival. *Journal of Applied Ecology* **47**, 1207–1215 (2010).
245. Bolnick, D. I. *et al.* The Ecology of Individuals: Incidence and Implications of Individual Specialization. *The American Naturalist* **161**, 1–28 (2003).
246. Smith, A. R., Graystock, P. & Hughes, W. O. H. Specialization on pollen or nectar in bumblebee foragers is not associated with ovary size, lipid reserves or sensory tuning. *PeerJ* **4**, e2599 (2016).

247. Chittka, L. & Thomson, J. D. Sensori-motor learning and its relevance for task specialization in bumble bees. *Behav Ecol Sociobiol* **41**, 385–398 (1997).

248. Maloof, J. E. The effects of a bumble bee nectar robber on plant reproductive success and pollinator behavior. *American Journal of Botany* **88**, 1960–1965 (2001).

249. Miller-Struttmann, N. E., Heise, D., Schul, J., Geib, J. C. & Galen, C. Flight of the bumble bee: Buzzes predict pollination services. *PLOS ONE* **12**, e0179273 (2017).

250. Ke, S.-R. *et al.* A Review on Video-Based Human Activity Recognition. *Computers* **2**, 88–131 (2013).

251. Wang, L. *et al.* Temporal Segment Networks: Towards Good Practices for Deep Action Recognition. in *ECCV* (2016). doi:10.1007/978-3-319-46484-8_2.

252. Zach, C., Pock, T. & Bischof, H. A Duality Based Approach for Realtime TV-L1 Optical Flow. in *Pattern Recognition* (eds. Hamprecht, F. A., Schnörr, C. & Jähne, B.) 214–223 (Springer, 2007). doi:10.1007/978-3-540-74936-3_22.

253. Ruiz, P. Understanding and visualizing ResNets. *Medium* https://towardsdatascience.com/understanding-and-visualizing-resnets-442284831be8 (2019).

254. Deng, J. *et al.* ImageNet: A large-scale hierarchical image database. in *2009 IEEE Conference on Computer Vision and Pattern Recognition* 248–255 (2009). doi:10.1109/CVPR.2009.5206848.

255. Wang, S. *et al.* denseflow. (2020).

# Appendix

## S1: Validation of Model Assumptions

### Do bees respond adaptively under variable floral reward conditions?

One major assumption of SimBee is that working memory enables bees to adaptively exploit flowers under highly variable reward conditions. To test this assumption, we randomly placed 4000 plants of each of four species (Sp1-4) in the environment with the following reward properties: reward mean and range for Sp 1=0.7(0.45,0.95); Sp 2=0.6 (0.35,0.85); Sp 3=0.5 (0.5,0.5), Sp 4=0.5 (0.0,1.0) (Scenario 1 condition, see Table 9). Individual bees were allowed to forage under these floral conditions for one sampling and foraging cycle (approximately 100 flower visits) either with the ability to remember reward values of the previous 0 flower visits (no memory condition) or 10 flower visits (memory condition). Figure 25a shows the proportion of visits to the plant species offering the greatest average reward (Sp1) for bees in each group. As expected, bees with memory quickly learned the identity of the plant species offering the greatest average reward and chose that species with high frequency compared with the baseline random choices of the bees without memory. Note that accuracy jumped when bees with memory entered foraging mode (at about 100 time units), and increased as bees filled their memory with more rewarding flowers. Because of the limited capacity of working memory and the variability of the floral rewards in the system, bees did not have 'complete' knowledge of floral reward properties of plant species and therefore did not exclusively visit the most rewarding species (Sp 1); in addition, Sp1 would not always be present in a bee's field-of-view. Memory also increased the amount of floral reward collected by bees over time (Figure 25b), demonstrating that memory is an adaptive trait in our system.

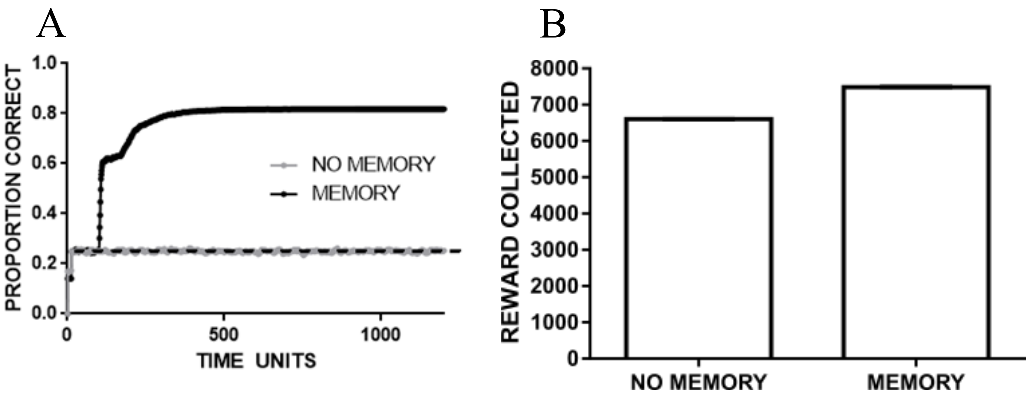

*Figure 25: Memory enables bees to make adaptive choices under variable floral reward conditions. Individual bees were allowed to choose among four plant species with different mean reward values over one sampling and foraging cycle. (a) Proportion of bee visits to the plant species with the greatest average reward. Initial visits were made with bees in sampling mode (not utilizing memory). Dotted line shows the expected proportion of visits to the most rewarding (correct) flower type assuming random choice. (b) Amount of reward collected by bees with the ability to remember reward values of the previous 0 (no memory) or 10 flower visits (memory). N=100 simulation runs for each memory condition. Mean+/- SD*

## Does plant reproductive success reflect bee foraging patterns?

Another major assumption of SimBee is that plant reproductive success depends on the type of pollen transferred among flowers (determined by flower choice behavior of bees) and amount of pollen transferred (determined by the number of bees in the system). To test these assumptions, we randomly placed 4000 plants of each of four species (Sp1-4) in the environment. Flowers of all species offered the same average reward (mean (range) = 0.58 (0.48, 0.68)). Each plant started the season with 6 empty slots that can either produce a seed (receive conspecific pollen) or be prevented from doing so in the future (receive heterospecific pollen; blocked slot) after each bee visit.

To test for effects of pollen limitation (fewer bees in the system) on seed production, all individuals were assigned a memory capacity of 1 and were allowed to forage in the following population sizes: 20, 50, 100, 200 individuals. Data were collected in foraging mode only. A memory capacity of 1 results in each individual bee always visiting only a single plant species. On average, all floral species were visited by approximately equal numbers of bees. Consequently, no slots were blocked for any species, and all seed production depended on the number of individual bees in the system. Results showed that seed production increased with bee population size (Figure 26a), indicating that plant reproductive success is limited by the amount of pollen transferred by bees. 100 bees were sufficient to almost completely fill all plant slots under these conditions.

To test the effects of flower choice behavior, bees were next assigned with one of the following memory capacities: remembered the previous 0 (no memory), 1, 5, 10, and 20 flower visits. Given the floral resource conditions, bees with a greater memory capacity are expected to switch between flowers of different plant species more frequently because they are better able to determine that average reward level is the same among plant species. Consequently, increased memory capacity is also expected to be associated with reduced seed production due to the associated increase in interspecific pollen transfer. Bees with 0 memory are expected to switch randomly among species, and thus should switch species for 75% of visits, providing a baseline for expected seed production. A bee population of 100 was used for each memory capacity. Data for bees and plants were generated in foraging mode only, over one season. We recorded the frequency of bee movements between flowers of the same species, or ‘like-like’ moves, which provides a direct measure of pollination efficiency when the number of bees is not limiting (i.e., seed production is based solely on the foraging decisions of bees). As expected, seed production increased with frequency of like-like moves by bees (Figure 26b), indicating that the foraging decisions of bees have a direct effect on plant fitness.

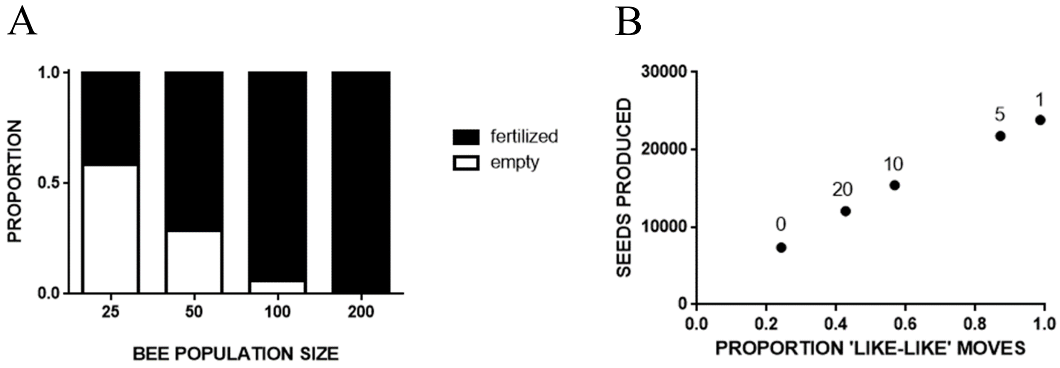

*Figure 26: Plant reproductive success depends on availability of bumblebee pollinators and their foraging choices. (a) Proportion of empty and fertilized (seed produced) seed slots at the end of the season as a function of bee population size. All bees had a memory capacity of 1 and therefore did not move pollen between different species; seed slots were always either empty or fertilized, and never blocked. (b) End of season seed production as a function of moves between flowers of the same species (like-like moves) for bees with different memory capacities. Bee population size was 100. Numbers above each point correspond to bee memory capacity. N=100 simulation runs for each experimental condition.*

## S2: Bumblebee-Inspired Vehicular Communication Algorithms

The challenge of adapting to changing and dynamic environments is not unique to pollinators. Modern vehicles also communicate with each other on the road using a limited range of wireless frequencies. As vehicle-to-vehicle (V2V) communication increases with vehicle density, spectrum scarcity issues are expected to emerge. A potential solution to this issue is to leverage underutilized wireless spectrum elsewhere, such as in the digital television spectrum band, using an approach called Vehicular Dynamic Spectrum Access (VDSA) [3], [4]. The fundamental idea behind DSA is to use unoccupied channels without interfering with the licensed users, i.e., primary users (PUs), of the frequency bands. However, channel occupancy and interference levels vary widely over time, creating a dynamic environment that is difficult to exploit. Channel energy – a combined measure of in-band and out-of-band interference – is a measure similar but inverse to nectar rewards in a floral environment. While the challenge of VDSA may be relatively novel in the electrical engineering field, there are numerous parallels between the wireless environment and the natural environment of pollinators such as bumblebees.

Models of bumblebee decision-making can provide a framework for VDSA based on the adaptive behavioral responses of pollinators in a complex and variable environment. In this context, the challenge of minimizing channel energy for each vehicle's communication contains many direct parallels to bumblebees maximizing their nectar intake while foraging. Figure 27 summarizes these comparisons and serves a foundation for the collaborative work with the Wyglinski lab at WPI. This collaboration resulted in the works listed below. These were written by Kuldeep Gill. Kevin Heath helped translate bumblebee algorithms into vehicular communications.

Gill, K. S. *et al.* Memory Matters: Bumblebee Behavioral Models for Vehicle-to-Vehicle Communications. *IEEE Access* **6**, 25437–25447 (2018).

Gill, K., Heath, K. N., Gegear, R. J., Ryder, E. F. & Wyglinski, A. M. On the Capacity Bounds for Bumblebee-Inspired Connected Vehicle Networks via Queuing Theory. in *2018 IEEE 87th Vehicular Technology Conference (VTC Spring)* 1–6 (2018). doi:10.1109/VTCSpring.2018.8417762.

Gill, K. S. *et al.* Experimental Test-Bed for Bumblebee-Inspired Channel Selection in an Ad-Hoc Network. in *2018 IEEE 88th Vehicular Technology Conference (VTC-Fall)* 1–5 (2018). doi:10.1109/VTCFall.2018.8690978.

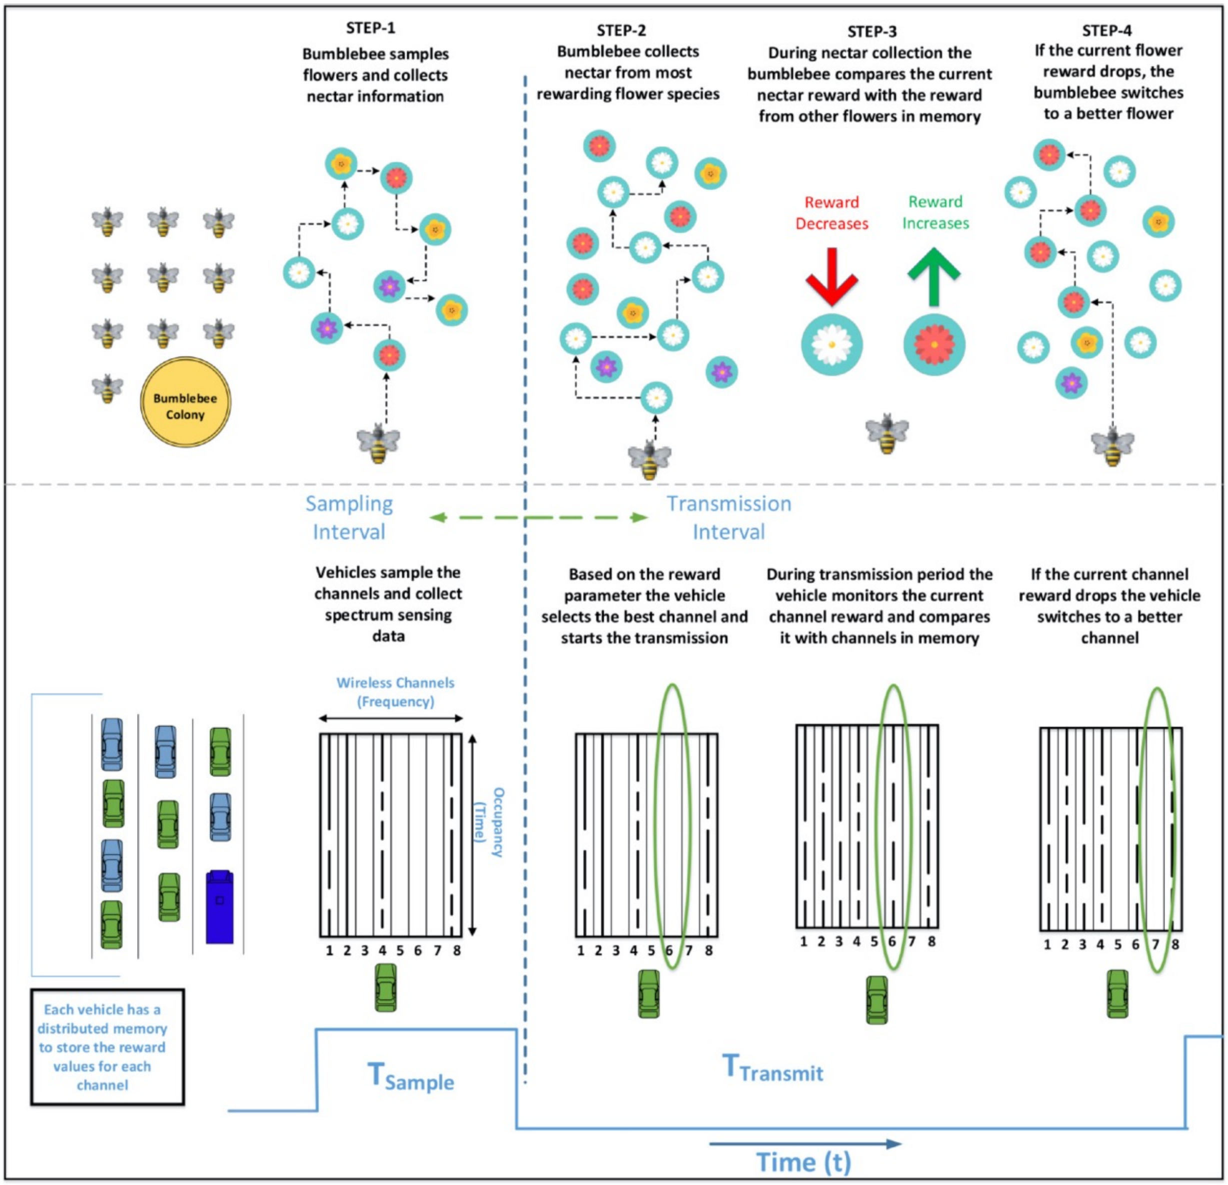

*Figure 27: The memory-based channel selection algorithm from bee (top) and vehicle (bottom) perspectives. Similar to each bumblebee, each vehicle is equipped with memory to store channel (floral) reward information, which is then used to select the channel (floral species) with the highest reward quality out of those available in sampling interval. During the transmission interval, the vehicles (bumblebees) use their current channels for communication (forage on current species) while simultaneously tracking the change in the reward level. The vehicles switch to a better channel (floral species) based on their memory if the current channel level drops to a lower value. Vehicles alternate between sampling ($T_{Sample}$) and transmission ($T_{Transmit}$) periods to track changes in a time-varying noisy resource environment.*

# S3: Simulation Results for Response to Change & Memory Dynamics in Bumblebees

## Scenario 2

*Table 24: Results for all decision-making strategies and memory models for the reward probability scenario.*

| fitness | memory size | memory model | memory alpha | softmax theta | env |
|---|---|---|---|---|---|
| 120.6 | NA | Win-Stay-Lose-Switch | NA | NA | 100_0 |
| 119.64 | 9 | exponential | 0.7 | NA | 100_0 |
| 119.64 | 1 | RLDelta | 0.25 | 50.5547 | 100_0 |
| 119.46 | 5 | ln | 0.25 | NA | 100_0 |
| 119.2 | 3 | Better than average | NA | NA | 100_0 |
| 118.98 | 4 | baseline | NA | NA | 100_0 |
| 118.64 | 3 | Tallying | NA | NA | 100_0 |
| 116.24 | 3 | Lexicographic | NA | NA | 100_0 |
| 60.04 | NA | Random | NA | NA | 100_0 |
| 92.4 | 1 | RLDelta | 0.65 | 99.4151 | 80_20 |
| 92.34 | 9 | exponential | 0.25 | NA | 80_20 |
| 91.56 | 6 | baseline | NA | NA | 80_20 |
| 91.4 | 5 | Lexicographic | NA | NA | 80_20 |
| 90.88 | 4 | Better than average | NA | NA | 80_20 |
| 90.1 | 6 | ln | 0.9 | NA | 80_20 |
| 88.98 | 3 | Tallying | NA | NA | 80_20 |
| 82.56 | NA | Win-Stay-Lose-Switch | NA | NA | 80_20 |
| 60.29 | NA | Random | NA | NA | 80_20 |
| 68.68 | 6 | Better than average | NA | NA | 60_40 |
| 68.62 | 10 | exponential | 0.05 | NA | 60_40 |
| 68.02 | 3 | baseline | NA | NA | 60_40 |
| 67.42 | 1 | RLDelta | 0.85 | 18.3484 | 60_40 |
| 66.9 | 9 | Lexicographic | NA | NA | 60_40 |
| 66.86 | 3 | Tallying | NA | NA | 60_40 |
| 66.04 | 9 | ln | 0.75 | NA | 60_40 |
| 64.86 | NA | Win-Stay-Lose-Switch | NA | NA | 60_40 |
| 60.45 | NA | Random | NA | NA | 60_40 |

## Scenario 3

*Table 25: Full results for the decision-making strategies and memory models for the frequency of change scenario.*

| fitness | memory size | memory model | memory alpha | softmax theta | env |
|---|---|---|---|---|---|
| 225.76 | 1 | RLDelta | 0.8 | 73.0256 | 60 visits |
| 222.12 | 3 | baseline | NA | NA | 60 visits |
| 221.2 | 3 | Better than average | NA | NA | 60 visits |
| 220.72 | 3 | Tallying | NA | NA | 60 visits |
| 219.2 | 10 | ln | 0.65 | NA | 60 visits |
| 217.72 | 9 | exponential | 0.75 | NA | 60 visits |
| 216 | 7 | Lexicographic | NA | NA | 60 visits |
| 205.67 | NA | Win-Stay-Lose-Switch | NA | NA | 60 visits |
| 148.97 | NA | Random | NA | NA | 60 visits |
| 218.32 | 1 | RLDelta | 0.85 | 25.6951 | 40 visits |
| 217.52 | 3 | exponential | 0.65 | NA | 40 visits |
| 217.44 | 3 | ln | 0.15 | NA | 40 visits |
| 211.84 | 3 | baseline | NA | NA | 40 visits |
| 208.48 | 3 | Tallying | NA | NA | 40 visits |
| 204.76 | 3 | Better than average | NA | NA | 40 visits |
| 203.32 | 5 | Lexicographic | NA | NA | 40 visits |
| 202.3 | NA | Win-Stay-Lose-Switch | NA | NA | 40 visits |
| 149.37 | NA | Random | NA | NA | 40 visits |
| 202.92 | 3 | ln | 0.7 | NA | 20 visits |
| 202.76 | 1 | RLDelta | 0.95 | 46.5936 | 20 visits |
| 200.33 | NA | Win-Stay-Lose-Switch | NA | NA | 20 visits |
| 192.72 | 3 | Lexicographic | NA | NA | 20 visits |
| 189.84 | 3 | exponential | 0.9 | NA | 20 visits |
| 185.96 | 3 | Better than average | NA | NA | 20 visits |
| 183.2 | 3 | Tallying | NA | NA | 20 visits |
| 170.05 | 4 | baseline | NA | NA | 20 visits |
| 150.27 | NA | Random | NA | NA | 20 visits |