

RESEARCH ARTICLE

Abstract concept learning in a simple neural network inspired by the insect brain

Alex J. Cope^{1,2*}, Eleni Vasilaki^{1,2}, Dorian Minors³, Chelsea Sabo^{1,2}, James A. R. Marshall^{1,2}, Andrew B. Barron³

1 Department of Computer Science, University of Sheffield, Sheffield, UK, **2** Sheffield Robotics, University of Sheffield, Sheffield, UK, **3** Department of Biological Sciences, Macquarie University, Sydney, Australia

* a.cope@sheffield.ac.uk



Abstract

The capacity to learn abstract concepts such as ‘sameness’ and ‘difference’ is considered a higher-order cognitive function, typically thought to be dependent on top-down neocortical processing. It is therefore surprising that honey bees apparently have this capacity. Here we report a model of the structures of the honey bee brain that can learn sameness and difference, as well as a range of complex and simple associative learning tasks. Our model is constrained by the known connections and properties of the mushroom body, including the protocerebral tract, and provides a good fit to the learning rates and performances of real bees in all tasks, including learning sameness and difference. The model proposes a novel mechanism for learning the abstract concepts of ‘sameness’ and ‘difference’ that is compatible with the insect brain, and is not dependent on top-down or executive control processing.

OPEN ACCESS

Citation: Cope AJ, Vasilaki E, Minors D, Sabo C, Marshall JAR, Barron AB (2018) Abstract concept learning in a simple neural network inspired by the insect brain. *PLoS Comput Biol* 14(9): e1006435. <https://doi.org/10.1371/journal.pcbi.1006435>

Editor: Joseph Ayers, Northeastern University, UNITED STATES

Received: April 24, 2018

Accepted: August 15, 2018

Published: September 17, 2018

Copyright: © 2018 Cope et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All files are available on GitHub at (<https://github.com/BrainsOnBoard/bee-concept-learning>).

Funding: JARM and EV acknowledge support from the Engineering and Physical Sciences Research Council (grant numbers EP/J019534/1 and EP/P006094/1), (<https://epsrc.ukri.org/>). JARM and ABB acknowledge support from a Royal Society International Exchanges Grant (<https://royalsociety.org/>). ABB is supported by an Australian Research Council Future Fellowship Grant 140100452 and Australian Research Council Discovery Project

Author summary

Is it necessary to have advanced neural mechanisms to learn abstract concepts such as sameness or difference? Such tasks are usually considered a higher order cognitive capacity, dependent on complex cognitive processes located in the mammalian neocortex. It has always been astonishing therefore that honey bees have been shown capable of learning sameness and difference, and other relational concepts. To explore how an animal like a bee might do this here we present a simple neural network model that is capable of learning sameness and difference and is constrained by the known neural systems of the insect brain, and that lacks any advanced neural mechanisms. The circuit model we propose was able to replicate bees’ performance in concept learning and a range of other associative learning tasks when tested in simulations. Our model proposes a revision of what is assumed necessary for learning abstract concepts. We caution against ranking cognitive abilities by anthropomorphic assumptions of their complexity and argue that by application of neural modelling, it can be shown that comparatively simple neural structures are sufficient to explain different cognitive capacities, and the range of animals that might be capable of them.

Grant DP150101172 (<http://www.arc.gov.au/>). The funders played no role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Introduction

Abstract concepts involve the relationships between things. Two simple and classic examples of abstract concepts are ‘sameness’ and ‘difference’. These categorise the relative similarity of things: they are properties of a relationship between objects, but they are independent of, and unrelated to, the features of the objects themselves. The capacity to identify and act on abstract relationships is a higher-order cognitive capacity, and one that is considered critical for any operation involving equivalence or general quantitative comparison [1–4]. The capacity to recognise abstract concepts such as sameness has even been considered to form the “very keel and backbone of our thinking” [5]. Several non-verbal animals have been shown to be able to recognise ‘sameness’ and ‘difference’ including, notably, the honey bee [6–9].

The ability of the honey bee to recognise ‘sameness’ and ‘difference’ is interesting, as the learning of abstract concepts is interpreted as a property of the mammalian neocortex, or of regions of the avian pallium [10–12] and to be a form of top-down executive modulation of lower-order learning mechanisms (complex neural processing affecting less complex processing, see [13] for more details) [4, 12]. This interpretation has been reinforced by the finding that activity of neurons in the prefrontal cortex of rhesus monkeys (*Macaca mulatta*) correlates with success in recognising sameness in tasks [11, 12]. The honey bee, however, has nothing like a prefrontal cortex in its much smaller brain.

In this paper we use a modelling approach to explore how an animal like a honey bee might be able to solve an abstract concept learning task. To consider this issue we must outline in more detail how learning of sameness and difference has been demonstrated in honey bees, and originally in other animals.

A family of ‘match-to-sample’ tasks has been developed to evaluate sameness and difference learning in non-verbal animals. In these tasks animals are shown a sample stimulus followed, after a delay, by two stimuli: one that matches the sample and one that does not. Sometimes delays of varying duration have been imposed between the presentation of the sample and matching stimuli to test duration of the ‘working memory’ required to perform the task [1, 14]. This working memory concept is likened to a neural scratchpad that can store a short term memory of a fixed number of items, previously seen but no longer present [15]. Tests in which animals are trained to choose matching stimuli are described as Match-to-Sample (MTS) or Delayed-Match-To-Sample (DMTS) tasks, and tests in which animals are trained to choose the non-matching stimulus are Not-Match-To-Sample (NMTS) or Delayed-Not-Match-To-Sample (DNMTS) tasks.

On their own, match-to-sample tasks are not sufficient to show concept learning of sameness or difference. For this it is necessary to show, having been trained to select matching or non-matching stimuli, that the animal can apply the concept of sameness or difference in a new context [4]. Typically this is done by training animals with one set of stimuli and testing whether they can perform the task with a new set of stimuli [6–9]; this is referred to as a transfer test.

In a landmark study [8] showed that honey bees can learn both sameness and difference. They could learn both DMTS and DNMTS tasks and generalise performance in both tasks to tests with new, previously unseen, stimuli [8]. In this study free-flying bees were trained and tested using a Y-maze in which the sample and matching stimuli were experienced sequentially during flight, with the sample at the entrance to the maze and the match stimuli at each of the y-maze arms. Bees could solve and generalise both DMTS and DNMTS tasks when trained with visual stimuli, and could even transfer the concept of sameness learned in an olfactory DMTS task to a visual DMTS task, showing cross-modal transfer of the learned concept of sameness [8]. Bees took 60 trials to learn these tasks [8]; this is much longer than learning a

simple olfactory or visual associative learning task, which can be learned by bees in 3 trials [16]. Their performance in DMTS and DNMTS was not perfect either; the population average for performance in test and transfer tests was around 75%, but they could clearly perform at better than chance levels [8] in both.

The concept of working memory is crucial for solving a DMTS/DNMTS task, as information about the sample stimulus is no longer available externally to the animal when choosing between the match stimuli. If there is no neural information that can identify the match then the task cannot be solved. We therefore must identify in the honeybee a candidate for providing this information in order to produce a model that can solve the task.

A previous model by [17] demonstrates DMTS and DNMTS with transfer, however the model contains many biologically unfounded mechanisms that are solely added for the purpose of solving these tasks, and the outcome of these additions disagrees with neurophysiological, and behavioural evidence. We instead take an approach of constraining our model strongly to established neurophysiology and neuronanatomy, and demonstrating behaviour that matches that of real bees. We will compare this model to the model presented here further in the Discussion. As the problem (without transfer) is a binary XOR, it can be solved by a network with at least one hidden layer, however this would not solve the generalisation aspect of the problem.

The honey bee brain is structured as discrete regions of neuropil (zones of synaptic contact). These are well described, as are the major tracts connecting them [18]. The learning pathways have been particularly intensely studied (*e.g.* [19–22]). The mushroom bodies (*corpora pedunculata*) receive processed olfactory, visual and mechanosensory input [23] and are a locus of multimodal associative learning in honey bees [19]. They are essential for reversal and configural learning [4, 24, 25]. Avarguès-Weber and Giurfa [4] have argued the mushroom bodies to be the most likely brain region supporting concept learning, because of their roles in stimulus identification, classification and elemental learning [19, 22, 26]. Yet it is not clear how mushroom bodies and associated structures might be able to learn abstract concepts that are independent of any of the specific features of learned stimuli and, crucially, how the identity of the sample stimulus could be represented. Solving such a problem requires two computational components. First, a means of storing the identity of the sample stimulus, a form of working memory; second, a mechanism that can learn to use this stored identity to influence the behaviour at the decision point. Below we propose a model of the circuitry of the honey bee mushroom bodies that can perform these computations and is able to solve DMTS and DNMTS tasks.

Results

Key model principles: A circuit model inspired by the honey bee mushroom bodies

We explored whether a neural circuit model, inspired and constrained by the known connections of the honey bee mushroom bodies, is capable of learning sameness and difference in a DMTS and DNMTS task (Fig 1). Full details of the models can be found in Methods. A sufficiency table for the key mechanisms and the following results is shown in Table 1.

The mushroom body has previously been modelled as an associative network consisting of three neural network layers [26, 27], comprised of input neurons (IN) providing processed olfactory, visual and mechanosensory inputs [23, 28], an expansive middle layer of Kenyon cells (KC) which enables sparse-coding of sensory information for effective stimulus classification [22], and finally mushroom body extrinsic neurons (EN) which output to premotor regions of the brain and can be considered (at this level of abstraction) to activate different

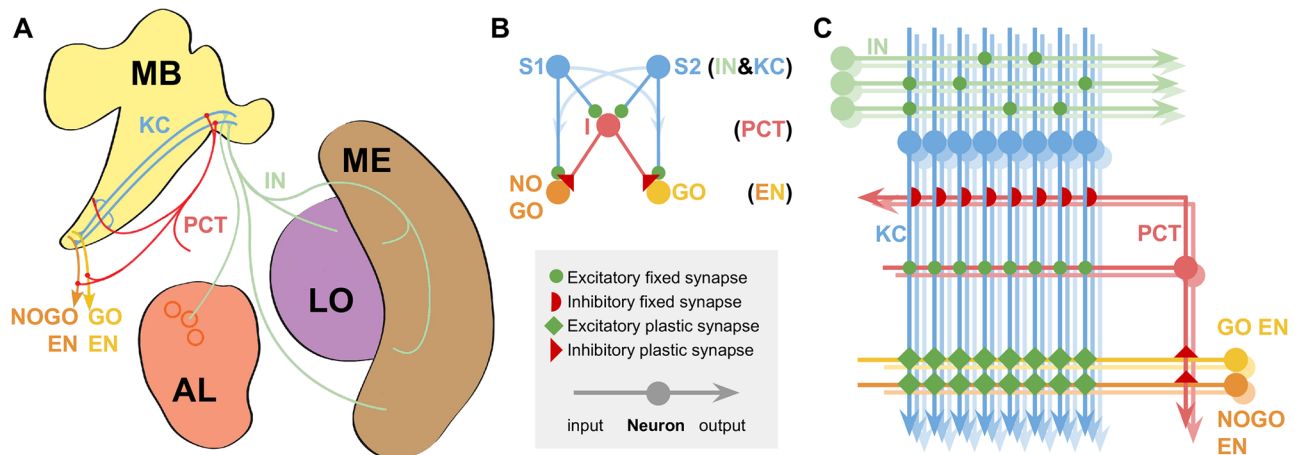


Fig 1. Models of the mushroom bodies based on known neuroanatomy. A Neuroanatomy: MB Mushroom Bodies; AL Antennal Lobe glomeruli (circles); ME & LO Medulla and Lobula optic neuropils. The relevant neural pathways are shown and labelled for comparison with the model. B Reduced model; neuron classes indicated at righthand side of sub-figure. C Full model, showing the model connectivity and indicating the approximate relative numbers of each neuron type. Colour coding and labels are preserved throughout all the diagrams for clarity. Excitatory and inhibitory connections indicated as in figure legend. Key of neuron types: KC, Kenyon Cells; PCT, Protocerebellar Tract neurons; IN, Input Neurons (olfactory or visual); EN, Extrinsic MB Neurons from the GO and NOGO subpopulations, where the subpopulation with the highest sum activity defines the behavioural choice in the experimental protocol (Fig 4).

<https://doi.org/10.1371/journal.pcbi.1006435.g001>

possible behavioural responses [22, 26]. Here, for simplicity, we consider the EN as simply two subpopulations controlling either ‘go’ or ‘no-go’ behavioural responses only, which allow choice between different options via sequential presentations where ‘go’ chooses the currently presented option. Connections between the KC output and ENs are modifiable by synaptic plasticity [26, 29–31] supporting learned changes in behavioural responses to stimuli.

As outlined in the introduction, we require two computational mechanisms for solving the DMTS/DNMTS task. First is a means of storing the identity of the sample stimulus. Second is learning to use this identity to drive behaviour and solve the task. Moreover this learning must generalise to other stimuli. The computational complexity of this problem should not be underestimated; either the means of storing the identity of the sample, or the behavioural learning, must generalise to other stimulus sets. The bees were not given any reward with the transfer stimuli in [8]’s study, so no post-training learning mechanism can explain the transfer performance. In addition, during the course of the experiment of [8] each of the two stimuli were used as the match, i.e. for stimuli A and B the stimulus at the maze entrance were alternated between A and B throughout the training phase of the experiment. This requires, therefore, that the bees have a sense of stimulus ‘novelty’, and can associate novelty with a behaviour: either approach for DNMTS, or avoid for DMTS. With one training set the

Table 1. Sufficiency table showing which learning mechanisms in the model are required for each result.

Result	KC-EN learning	PCT-EN learning
DMTS	N	Y
DNMTS	N	Y
Single PER	Y	N
+ve pattern PER	Y	N
-ve pattern PER	Y	N
Differential	Y	N

<https://doi.org/10.1371/journal.pcbi.1006435.t001>

problem is solvable as delayed paired non-elemental learning tasks (e.g. responding to stimuli A and B together, but not individually), however with the transfer of learning to new stimulus sets such an approach does not solve the whole task.

There is one feature of the Kenyon Cells which can fulfill this computational requirement for novelty detection, that of sensory accommodation. In honey bees, even in the absence of reward or punishment, the KC show a stark decrease in activity between initial and repeated stimulus presentations of up to 50%, an effect that persists over several minutes [32]. This effect is also found in *Drosophila melanogaster* [33], where there is additionally a set of mushroom body output neurons that show even starker decreases in response to repeated stimuli, and which respond to many stimuli with stimulus specific decreases (thus making them clear novelty detectors), however such a neuron has not been found in bees to date. This response decrease in Kenyon Cells found in flies and bees is sufficient to influence behaviour during a trial but, given the decay time of this effect, not likely to influence subsequent trials. The mechanism behind this accommodation property is not known, and therefore we are only able to model phenomenologically, which we do by reducing the strength of the KC synapses for the sample stimulus by a fixed factor, tuned to reproduce the reduction in total KC output found by Szyszka et al. [32] (see Fig 2 panel E). However it should be noted that stimulus-specific adaptation is shown in many species and brain areas, and can be explained by short-term plasticity mechanisms [34–36] and architectural constraints only; see for instance [37]. Exploiting the faculty of the units to adapt to repeated stimuli allows the network to distinguish between repeated stimuli (“same”) and non-repeated, i.e. novel “different”) stimuli.

Having identified our first computational mechanism, a memory trace in the form of reduced KC output for the repeated stimulus, we need only to identify the second, a learning mechanism that can use this reduced KC output to drive behaviour to choose the correct (matching or non-matching) arm of the y-maze. If this learning mechanism exists at the output synapses of the KC it is either specific to the stimulus—if using a pre-postsynaptic learning rule—and therefore cannot transfer, or it utilises a postsynaptic-only learning rule. Initially the postsynaptic learning rule appears a plausible solution, however we must consider that bees can learn both DMTS and DNMTS, and that learning can only occur when the bee chooses to ‘go’. This creates a contradiction, as postsynaptic learning will proportionally raise both the weaker (repeated) stimulus activity, as well as the stronger (non-repeated) stimulus activity in the GO EN subpopulation. To select ‘go’ the GO activity for the currently presented stimulus must be larger than the activity in the NOGO subpopulation, which is fixed. Therefore we face the contradiction that in the DMTS case the weaker stimulus response must be higher than the stronger one in the GO subpopulation with respect to their NOGO subpopulation counterpart responses, yet in the DNMTS case the converse must apply. No single postsynaptic learning rule can fulfil this requirement.

A separate set of neurons that can act as a relay between the KC and behaviour is therefore required to solve both DMTS and DNMTS tasks. A plausible candidate is the inhibitory neurons that form the protocerebellar tract (PCT). These neurons have been implicated in both non-elemental olfactory learning [25] and regulatory processes at the KC input regions. They also project to the KC output regions [38–41], where there are reward-linked neuromodulators and learning-related changes [20, 42]. These neurons are few in number in comparison to the KC population, and some take input from large numbers of KC [43]. We therefore propose that, in addition to their posited role in modulating and regulating the input to the KC based on overall KC activity [43], these neurons could also regulate and modulate the activity of the EN populations at the KC output regions. Such a role would allow, via synaptic plasticity, a single summation of activity from the KCs to differentially affect both their inputs and outputs. If we assume a high threshold for activity for the PCT neurons (again consistent with their

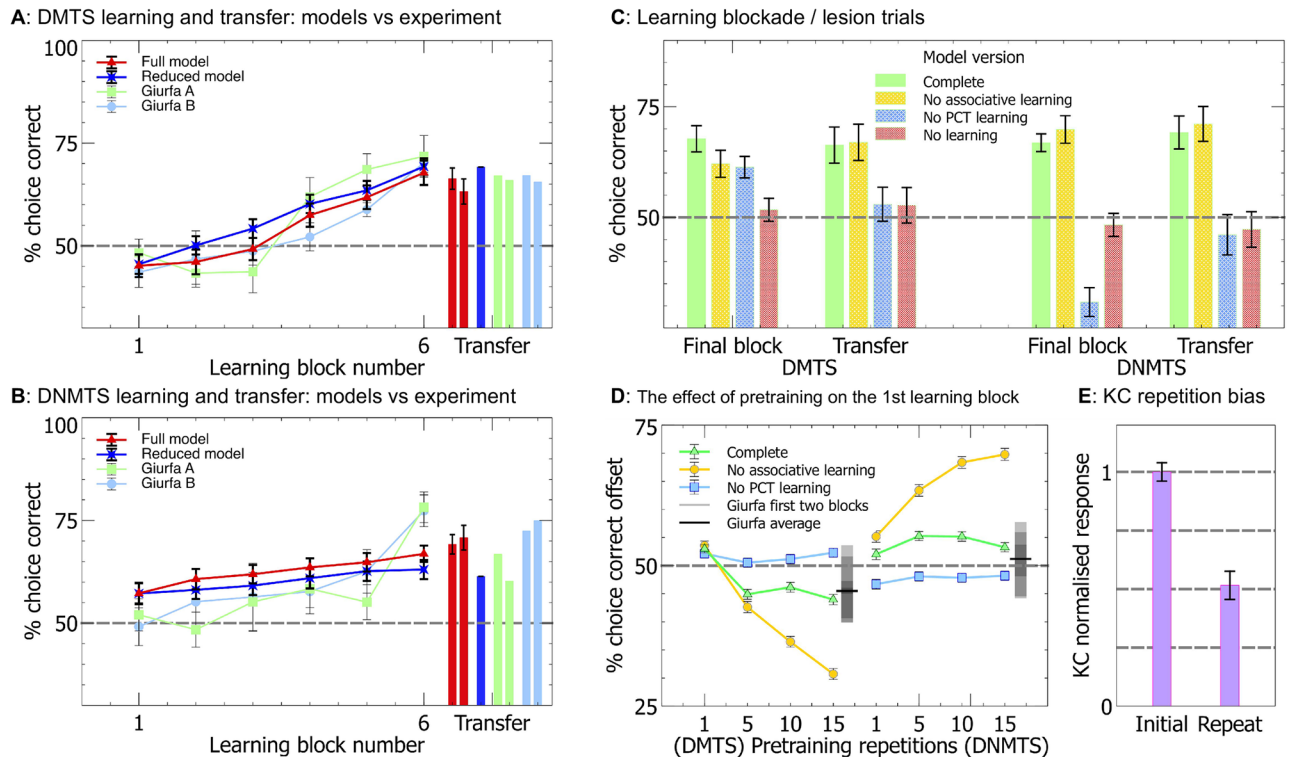


Fig 2. The full and reduced versions of our model reproduce the transfer of sameness and difference learning. A & B The average percentage of correct choices made by the model and real bees within blocks of ten trials as the task is learned (lines), along with the transfer of learning onto novel stimulus sets (bars). Both versions of the model reproduce the pattern of learning acquisition for DMTS (Full: N = 338, Reduced: N = 360) and DNMTS (Full & Reduced: N = 360) found when testing real bees (test for learning: $P < 0.0001$), along with the transfer of learning ($P < 0.0001$). For DMTS *Giurfa A* & *B* are the data from Experiments 1 & 2 respectively from *Giurfa et al.* [8], and for DNMTS *Giurfa A* & *B* are the data from Experiments 5 & 6 respectively from the same source. For an explanation of the initial offsets from chance for the model please see the text for panel D. C The blockade of plasticity from the MB and PCT pathways shows that the PCT pathway is necessary and sufficient for sameness and difference learning in the full model. All non-overlapping SEM error bars are significantly different. D PCT pathway learning in the absence of associative learning leads to preference for non-matching stimuli following pre-training, demonstrating that learning in the associative pathway changes the form of the sameness and difference acquisition curves. The equivalent offsets and error ranges for the first two blocks of *Giurfa* Experiments 1, 2, 5 & 6 along with the averages for DMTS and DNMTS for these blocks are shown alongside the model data for comparison as overlapping grey boxes—overlapping boxes create darker regions, thus the area of greatest darkness is the point where the most of the error ranges overlap. E The average activity of the model KC neurons when presented with repeated stimuli.

<https://doi.org/10.1371/journal.pcbi.1006435.g002>

proposed role) such that repeated stimuli would not activate the PCT neurons but non-repeated stimuli would, it is then possible for synaptic plasticity from the PCT neurons to the EN to solve the DMTS and DNMTS tasks and, vitally, transfer that learning to novel stimuli. We do not propose that this is the purpose of these neurons, but instead that it is a consequence of their regulatory role.

We present two models inspired by the anatomy and properties of the honey bee brain that are computationally capable of learning in DMTS and DNMTS tasks, and the generalisation of this learning to novel stimuli (Fig 1).

Our first, reduced, model is a simple demonstration that the key principles outlined above can solve DMTS and DNMTS tasks, and generalise the learning to novel stimulus sets. By simplifying the model in this way the computational principles are readily apparent. Such a simple model, however, cannot demonstrate that associative learning in the KC to EN synapses does not interfere with learning in the PCT to EN synapses or vice versa. For this we present a full model that includes the associative learning pathway from the KC to the EN, and demonstrate

that this model can not only solve DMTS and DNMTS with transfer to novel stimuli, but can also solve a suite of associative learning tasks in which the MB have been implicated. The results of computational experiments performed with these models are presented below. The full model addresses the interaction of the PCT to EN learning and the KC to EN learning, as well as suggesting a possible computational role of the PCT to EN synaptic pathway in regulating the behavioural choices driven by the MB output, which we present in the Discussion.

A reduced model of the core computational principles produces sameness and difference learning, and transfers this learning to novel stimuli

The reduced model is shown in Fig 1 Panel B and model equations are presented in Methods. The input nodes S_1 and S_2 represent the two alternative stimuli, where we have reduced the sparse KC representation into two non-overlapping single nodes for simplicity, and as such we do not need to model the IN input neurons separately. Node I (which corresponds to the PCT neurons, again reduced to a single node for simplicity) represents the inhibitory input to the output neurons GO and NOGO. Nodes S_1 and S_2 project to nodes I and to GO and NOGO with fixed excitatory weighted connection. Finally, node I projects to GO and NOGO with plastic inhibitory weighted connections. Node I is thresholded so that it only responds to novel stimuli.

Fig 2 panels A and B show the performance of the reduced model bees (as well as the full model) for task learning and transfer to novel input stimuli, alongside experimental data from [8]. While the reduced model solves the transfer of sameness and difference learning the pre-training process strongly biases the model towards non-repeated stimuli, proportional to the number of pretraining trials. Notably, this bias in the reduced model is different to that found in the full model, which we discuss below.

The model operates by adjusting the weights between the I and the GO to change the likelihood of choosing the non-matching stimulus. Since only connections from the I (representing the PCT neuron) to GO neurons are changed, the I to GO/NOGO weights are initialised to half the maximum weight value. Note that the I node is only active for the non-repeated stimulus, and this pathway has no effect for repeated stimuli. This means that if the weights are increased then non-repeated stimuli will have greater inhibition to the GO neuron, and therefore be less likely to be chosen. If the weights are decreased then non-repeated stimuli will have less inhibition to the GO neuron and therefore will be more likely to be chosen. As the conditions for changing the weights are only met when the non-repeated stimulus is chosen for ‘go’, this means that the model only learns on unsuccessful trial for DMTS (increasing the weight), or successful trials for DNMTS (decreasing the weight).

A full model is capable of sameness and difference learning, and transfers this learning to novel stimuli

The full model is shown in Fig 1 Panel C and model equations are presented in Methods. Fig 2 panel D shows the performance of the full model for the first block of learning following from different numbers of pretraining repetitions. When only the PCT pathway is plastic there is a large bias towards the non-repeated stimulus due to the pretraining, as found in the reduced model. This bias is reduced by the presence of the associative learning pathway, and the bias is independent of the number of pretraining trials for more than 5 trials. It should be noted that the experimental data [8] show indications of such a bias, in line with the results from the full model. The reduced model therefore requires fewer pretraining trials than the full model to produce a similar bias, which leads to the reduced model having large maladaptive behavioural biases for non-repeated stimuli if all stimuli are rewarded. This is important, as it suggests a

role for the PCT pathway in modulating the behavioural choice of the bee. This possible role is explored further in the Discussion.

Fig 2 panels A and B show the performance of the model bees compared with the performance of real bees from [8], and the reduced model. In both cases the trends found in the performance of the model bees match the trends found in the real bees for both task learning and transfer to novel stimuli. It is important to note the different forms of the learning in the DMTS and DNMTS paradigms, with DNMTS slower to learn. This is a direct consequence of the inhibitory nature of the PCT neurons; excitatory neurons performing the role of the PCT neurons in the model would lead to a reversal of this feature, with DMTS learning more slowly.

Learning in the PCT pathway of the full model is essential for transfer of learning to novel stimuli

We next sought to confirm that learning in the PCT neuron to EN pathway enabled generalisable learning of sameness and difference. Computational modelling provides powerful tools with which to do this, by comparing model performance when different elements are suppressed with the full model. We selectively suppressed the KC associative learning pathway, the PCT pathway learning, and all learning in the model. When a learning mechanism is suppressed this means that the synaptic weights stay the same throughout the training, but the pathway is otherwise active.

The results are summarised in Fig 2 panel C. It can clearly be seen that within our model learning in the PCT pathway is necessary for transfer of the sameness and difference learning to novel stimuli. Associative learning via the KC pathway alone has no effect on the transfer task performance compared to the fully learning-suppressed model. Unsuppressed associative learning leads to a preference for the matched stimulus, which has weaker KC activity, but this learning is specific to the trained stimuli, and does not transfer to novel stimuli.

Validation: The full model is capable of performing a range of conditioning tasks

Many models have reproduced the input neuron to Kenyon Cell to Extrinsic neuron pathway [26, 27, 44, 45], and these models demonstrate many forms of elementary and complex associative learning that have been attributed to the mushroom bodies. It is therefore important to demonstrate that in our model the PCT neuron pathway does not affect the reproduction of such learning behaviours. We therefore tested elemental and non-elemental associative learning undertaken by conditioning the Proboscis Extension Reflex (PER) in restrained bees, and reversal learning in free flying bees, as described in Methods. Our model is capable of reproducing the results found in experiments involving real bees, with the model's acquisition curves showing similar to the performance to the real bees. The results are shown in Fig 3, and details of the experimental paradigms used can be found in Methods.

Discussion

We have presented a simple neural model that is capable of learning the concepts of sameness and difference in Delayed Match to Sample (DMTS) and Delayed Not Match to Sample (DNMTS) tasks. Our model is inspired by the known neurobiology of the honey bee, and is capable of reproducing the performance of honey bees in a simulation of DMTS and DNMTS tasks. Our model therefore proposes a hypothesis for how animals like the honey bee might be able apparently to learn abstract concepts.

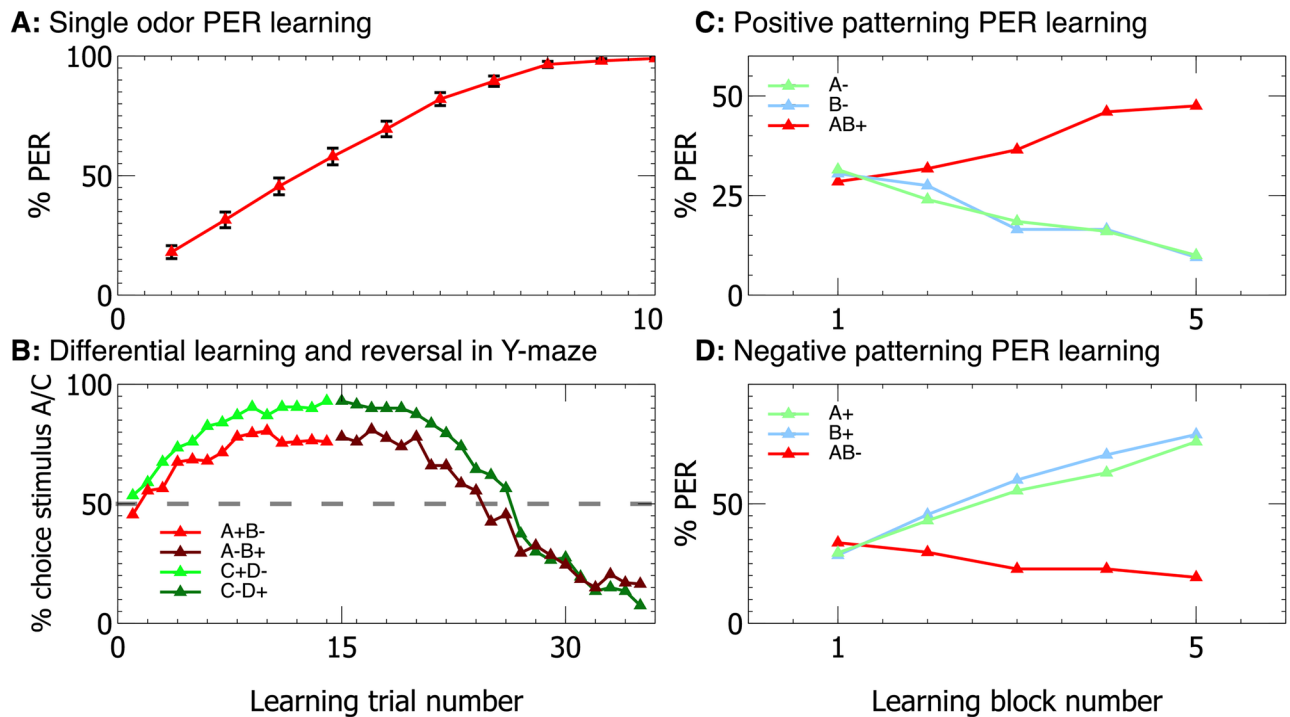


Fig 3. The full model is capable of performing a range of conditioning tasks. With modification of only the experimental protocol, our full model can successfully perform a range of conditioning tasks which can be performed by restrained (using the Proboscis Extension Reflex (PER) paradigm) and free flying bees. Performance closely matches experimental data with real bees (e.g. A: [46], B: [47], C & D: [48]).

<https://doi.org/10.1371/journal.pcbi.1006435.g003>

Abstract concept learning is typically described as a higher-order cognitive capacity [1, 4], and one that is dependent on a top-down modulation of simpler learning mechanisms by a higher cognitive process [49]. By contrast our model proposes a solution to sameness and difference learning in DMTS-style tasks with no top-down structure (that is, no complex mechanisms imposing themselves on simpler processes). The actions of the PCT neurons are integrated with the KC learning pathway and provide a parallel processing pathway sensitive to stimulus magnitude, rather than a top-down imposition of a learned concept of sameness or difference (Fig 1). This is a radical new proposal for how an abstract concept might be learned by an animal brain.

The first question we must ask when constructing a model regards plausibility. Our model (Fig 1) shows a close match to the neuroanatomical data for the mushroom bodies. Several computational requirements of our model match with experimental data, notably the sensory accommodation in the response of the KC neurons. Previous neural models based on this structure have proposed mechanisms for various forms of associative learning, including extinction of learning, and positive and negative patterning [17, 26, 45]. Our model is also capable of solving a range of stimulus-specific learning tasks, including patterning (Fig 3). No plausible previous model of the MB or the insect brain has been capable of learning abstract concepts, however.

As mentioned in the Introduction, a previous model by [17] demonstrates DMTS and DNMTS with transfer. Their motivation is the creation of a model for robotic implementation, rather than reproduction of behavioural observations from honey bees. While we suggest a role for the PCT neurons given experimental evidence of changes in the response of Kenyon Cells to repeated stimuli, Arena et al.'s model assumes resonance between brain regions that

is dependent upon the time after stimulus onset and the addition of specific neurons for 'Match' and 'Non-match'; there is no biological evidence for either of these assumptions. Furthermore, the outcome of these additions is an increase in Kenyon cell firing in response to repeated stimuli; this is in opposition to neurophysiological evidence from multiple insect species, including honey bees [32, 33]. In addition, Arena et al's proposed mechanism does not replicate the difficulty honey bees have in learning DMTS/DNMTS tasks, exhibiting learning in three trials, as opposed to 60 in real bees. In contrast, our model captures the rate and form of the learning found in real honey bees.

To enable a capacity for learning the stimulus-independent abstract concept of sameness or difference our model uniquely includes two interacting pathways. The KC pathway of the mushroom bodies retains stimulus-specific information and supports stimulus-dependent learning. The PCT pathway responds to summed activity across the KC population and is therefore largely independent of any stimulus-specific information. This allows information on stimulus magnitude, independent of stimulus specifics, to influence learning. Including a sensory accommodation property to the KCs [32] makes summed activity in the KCs in response to a stimulus sensitive to repetition, and therefore stimuli encountered successively (same) cause a different magnitude of KC response to novel stimuli (different) irrespective of stimulus specifics. This model is capable of learning sameness and difference rules in a simulation of the Y-maze DMTS and DNMTS tasks applied to honey bees (Fig 2), but in theory it could also solve other abstract concepts related to stimulus magnitude such as quantitative comparisons [4, 50].

Our model demonstrates a bias towards non-repeated stimuli, induced by the combination of sensory accommodation in the KC neurons and PCT learning during the pretraining phase, and largely mitigated by associative learning in the KC to EN synapses. This bias (see Fig 2) is indicated in the data from [8], and could be confirmed by further experimentation.

We note, however, that our model only supports a rather limited form of concept learning of sameness and difference. Learning in the model is dependent on sensory accommodation of the KCs to repeated stimuli [32]. This effect is transient, and hence the capacity to learn sameness or difference will be limited to situations with a relatively short delay between sample and matching stimuli. This limitation holds for honey bee learning of DMTS tasks [51], but many higher vertebrates do not have this limitation [52]. For example, in capuchins learning of sameness and difference is independent of time between sample and match [53]. We would expect that for animals with larger brains and a developed neocortex (or equivalent) many other neural mechanisms are likely to be at play to reinforce and enhance concept learning, enabling performance that exceeds that demonstrated for honey bees. Monkey pre-frontal cortex (PFC) neurons demonstrate considerable stimulus-specificity in matching tasks, and different regions appear to have different roles in coding the salience of these stimuli [54, 55]. Recurrent neural activity between these selective PFC neurons and lower-order neural mechanisms could support such time independence. Language-trained primates did particularly well on complex identity matching tasks and the ability to form a language-related mental representation of a concept might be the reason [56–58].

Wright and Katz [1] have utilised a more elaborate form of a MTS task in which vertebrates simultaneously learn to respond to sameness and difference, and are trained with large sets of stimuli rather than just two. They argue this gives less-ambiguous evidence of true concept learning since both sameness and difference are learned during training, and the large size of the training stimulus set encourages true generalisation of the concept. In theory our model could also solve this form of task, but it is unlikely a honey bee could. Capuchins, rhesus and pigeons required hundreds of learning trials to learn and generalise the sameness and difference concepts [1]. Bees would not live long enough to complete this training,

Finally as a consequence of our model we question whether it is necessary to consider abstract concept learning to be a higher cognitive process. Mechanisms necessary to support it may not be much more complex than those needed for simple associative learning. This is important because many behavioural scientists still adhere to something like Lloyd Morgan’s Canon [59], which proposes that “in no case is an animal activity to be interpreted in terms of higher psychological processes if it can be fairly interpreted in terms of processes which stand lower in the scale of psychological evolution and development” ([59] p59). Yet the Canon is therefore reliant on an unambiguous stratification of cognitive processes according to evolutionary history and development [60]. If abstract concept learning is in fact developmentally quite simple, evolutionarily old and phylogenetically widespread, then Morgan’s Canon would simply beg the question of why even more animals do not have this capacity [61]. We argue that far more information on the precise neural mechanisms of different cognitive processes, and the distribution of cognitive abilities across animal groups, is needed in order to properly rank capacities as higher or lower. In other words, so-called higher cognitive processes may be solved by relatively simple structures.

Methods

Model parameter selection

Many of the parameters of the model were fixed by the neuroanatomy of the honey bee, as well as the previous values and procedures described in [26], with the following modifications.

First, we increased the sparseness of the connectivity from the PN to the KC.

Second, the reduction in the magnitude of the KC output to repeated stimuli was tuned to replicate the magnitude of reduction described in [32].

Third, the learning rates were set so that acquisition of a single stimulus is rapid. In addition there are two ratios from this initial value that must be set. These are the ratio of the speed of excitatory associative learning in the Kenyon Cell to Extrinsic Neuron pathway to the inhibitory learning in the Protocerebellar Tract to Extrinsic Neuron pathway, and the ratio of the speed of acquisition when rewarded to the speed to extinction when no reward is given. We conservatively set both of these ratios to 2:1, with excitatory learning faster than inhibitory learning, and extinction faster than acquisition. The learning parameters used in the reduced model are used in the fitting of the full model.

Finally, we tuned the threshold value for the PCT neurons so that they only responded to a new stimulus, and not a repeated one.

A full list of the parameters can be found in Table 2.

Table 2. Model parameters; all parameters are in arbitrary units.

Name	Value	Name	Value
Full			
N_{IN}	144	N_{KC}	5000
N_{EN}	8	N_{PCT}	6
$p_{IN - >KC}$	0.02		
b	1.2	$b_s (l > 0)$	150
$b_s (l = 0)$	120		
Reduced			
c	80	d_0	1
Shared			
λ_e	0.06	λ_i	0.03
R_b	2/3		

<https://doi.org/10.1371/journal.pcbi.1006435.t002>

Reduced model

The reduced model is shown in Fig 1, and described in the text in Results. Here are the equations governing the model.

The input node S_1 projects to node I via a fixed excitatory weight of 1.0 and to GO and NOGO with excitatory weights w_{GO,S_1}^e and w_{NOGO,S_1}^e correspondingly (superscript denotes *excitatory* and subscript the connection from neuron S_1 to GO/NOGO). Similarly, node S_2 projects to I via an excitatory weight w_{I,S_2}^e and to GO and NOGO with excitatory weights w_{GO,S_2}^e and w_{NOGO,S_2}^e . Finally, node I projects to GO and NOGO with inhibitory weights $w_{GO,I}^i$ and $w_{NOGO,I}^i$ correspondingly. Node I is a threshold linear neuron with a cut-off at high values of activity x_{max} . Nodes GO and NOGO are linear neurons, with activities restricted between zero and one.

The model is described by the following equations, where only one input node S_1 or S_2 are active (but not both, as the bee observes one option at a time), where the activities of neurons I, GO and NOGO are calculated by:

$$I = S_i \Theta(S_i > \theta) \tag{1}$$

$$GO = w_{GO,S_i}^e S_i - w_{GO,I}^i I \tag{2}$$

$$NOGO = w_{NoGo,S_i}^e S_i - w_{NoGo,I}^i I, \tag{3}$$

where $i = 1, 2$ is an index taking values depending on which stimulus is present (S_1 or S_2), and neuronal activities of I, GO and NOGO are constrained between x_{min} and x_{max} . If a stimulus has been shown twice, during its second presentation there is a suppression of the neuronal activity that represents the specific stimulus, consistent with experimental findings [32]. This is modelled as a reduction by a factor of 0.7 of the value S_i for the repeated stimuli.

To calculate the probability of the behavioral outcome of GO or NOGO being the winner we use the following equation:

$$P(GO) = \frac{1}{1 + e^{-(c-d)(GO-NOGO)}}, \tag{4}$$

$$P(NOGO) = 1 - P(GO). \tag{5}$$

where c is a fixed coefficient and d a bias that increases linearly with the time it takes to make a decision, in the following way:

$$d = \frac{k}{d_o}, \tag{6}$$

with k being the number of consequent iterations for which GO has not been selected, set at zero at the beginning of each stimulus presentation. The parameter d_o is a constant, and selected so that the factor $c - d$ will always be positive. This parameterisation makes sure that the longer it takes for a decision GO to be made, the higher the probability that GO will be chosen at the next step. The full model explicitly implements this strategy iteratively, and validation of this approach can be seen in the close agreement between the behavioural outcomes of the two models.

Inhibitory synaptic weights w^i are learned using the following equation:

$$\Delta w^i = -\lambda_i (R - R_b) \text{ presynaptic activity} \times \text{postsynaptic activity}, \tag{7}$$

where λ_i is the learning rate of the inhibitory weights, reward $R = 1$ if reward is given, and zero in all other cases, R_b is a reward baseline, and the presynaptic (postsynaptic) activity is 1 if the

presynaptic (postsynaptic) neuron is active and 0 elsewhere. This is a reward-modulated Hebbian rule also known as a three factor rule [62, 63].

Additional neuronal inputs with similar connectivity as S_1 and S_2 , not shown explicitly in the diagram, are also present in the model simulations, and constructing the equations for these simply requires substitution of S_i for T_i in Eqs 1, 2 and 3. These represent the transfer stimuli and can be used following training to demonstrate transfer of learning. Details of training the model can be found in the Experiment subsection of the Methods.

Full model

The full model is shown in Fig 1. Our model builds on a well established abstraction of the mushroom body circuit (see [26, 27, 44]) to model simple learning tasks.

The main structure of the model consists of an associative network with three neural network layers. Adapting terminology and features from the insect brain we label these: input neurons (IN) (corresponding to S in the Reduced Model), a large middle layer of Kenyon cells (KC) (corresponding to the S to GO / NOGO connections in the Reduced Model), and a small output population of mushroom body extrinsic neurons (EN) separated into GO and NOGO subsets (as in the Reduced Model). The connections, c_{ij} , between the IN and KC are fixed, and are randomly selected from the complete connection matrix with a fixed probability $p_{IN \rightarrow KC} = 0.02$. Connections from the KC to the EN are plastic, consisting of a fully connected matrix. The connection strength between the j th KC and the k th EN (w_{jk}) can take a value between zero and one. The neural description used in the entire model is linear with a bottom threshold, and contains no dynamics, consisting of a summation over the inputs followed by thresholding at a value b via a Heaviside function Θ , with a linear response above the threshold value. This gives the associative model as the following, where the outputs of i th, j th and k th neurons of the IN, KC and EN populations are x_i , y_j and z_k respectively, and M describes the modulation of KC activity for the stimulus seen at the maze entrance:

$$M = \begin{cases} 1.0 & \text{: at maze entrance} \\ 1.0 & \text{: in arms if } y_j = 0 \text{ (KC inactive) at entrance} \\ 0.7 & \text{: in arms if } y_j > 0 \text{ (KC active) at entrance} \end{cases} \quad (8)$$

$$y_j = M \Theta \left(\sum_{i=0}^{N_{IN}} c_{ij} x_i - b \right) \left(\sum_{i=0}^{N_{IN}} c_{ij} x_i - b \right) \quad (9)$$

$$z_k = \Theta \left(\sum_{j=0}^{N_{KC}} w_{jk} y_j \right) \left(\sum_{j=0}^{N_{KC}} w_{jk} y_j \right) \quad (10)$$

where N_{IN} is the number of IN and N_{KC} is the number of KC.

DMTS generalisation is performed by the inhibitory protocerebral tract (PCT) neurons s_l (corresponding to I in the Reduced Model) described by the following equations:

$$s_l = \Theta \left(\sum_{j=0}^{N_{KC}} x_j - b_s \right) \left(\sum_{j=0}^{N_{KC}} x_j - b_s \right) \quad (11)$$

$$y_j = \Theta \left(\sum_{i=0}^{N_{IN}} c_{ij} x_i - b - \sum_{l=0}^{N_{PCT}} s_l^* \right) \left(\sum_{i=0}^{N_{IN}} c_{ij} x_i - b - \sum_{l=0}^{N_{PCT}} s_l^* \right) \quad (12)$$

$$z_k = \Theta \left(\sum_{j=0}^{N_{KC}} w_{jk} x_j - 0.5 \sum_{l=0}^{N_{PCT}} w_{lk}^i s_l \right) \left(\sum_{j=0}^{N_{KC}} w_{jk} x_j - 0.5 \sum_{l=0}^{N_{PCT}} w_{lk}^j s_l \right) \quad (13)$$

Where w_{lk} are the inhibitory weights between the *PCT* neurons. The $*$ denotes 10 iteration delayed activity from the *PCT* neurons due to delays in the *KC*->*PCT*->*KC* loop.

Learning takes place according to Eq (7), and the following equation for excitatory synapses:

$$\Delta w^e = \lambda_e (R - R_b) \text{ presynaptic activity} \times \text{postsynaptic activity}, \quad (14)$$

where λ_e is the learning rate of the excitatory weights, reward $R = 1$ if reward is given, and zero in all other cases, R_b is a reward baseline, and the presynaptic (postsynaptic) activity is 1 if the presynaptic (postsynaptic) neuron is active and 0 elsewhere.

Similarly to the reduced model, a decision is made when the *GO* EN subpopulation activity is greater than the *NOGO* EN population by a bias Rd , where d increases every time a *NOGO* decision is made by 10.0, and R is a uniform random number in the range $[-0.5, 0.5]$. To prevent early decisions the sum of the whole EN population activity must be greater than 0.1.

Experiment

Our challenge is to reproduce Giurfa et al's data demonstrating bees solving DMTS and DNMTS tasks [8]. To aid exploration of our model we simplify the task it must face, while retaining the key elements of the problem as faced by the honeybee. We therefore embody our model in a world described by a state machine. This simple world sidesteps several navigation problems associated with the real world, however we believe that for the sufficiency proof we present here such simplifications are acceptable—the ability of the honeybee to form distinct and consistent neural representations of the training set stimuli as it flies through the maze is a prerequisite of solving the task, and is therefore assumed.

The experimental paradigm for our Y-maze task is shown in Fig 2. The model bee is moved between a set of states which describe different locations in the Y-maze apparatus: at the entrance, in the central chamber facing the left arm, in the central chamber facing the right arm, in the left arm, in the right arm. When at the entrance or in the main chamber the bee is presented with a sensory input corresponding to one of the test stimuli. We can then set the test stimuli presented to match the requirements of a given trial (e.g. entrance (A), main chamber left (A), main chamber right (B) for DMTS when rewarding the left arm, or DNMTS when rewarding the right arm).

Experimental environment

The experimental environment consists of a simplified Y-maze (see Fig 2: main paper), in which the model bee can assume one of three positions: at the entrance to the Y-maze; at the choice point in front of the left arm; at the choice point in front of the right arm. At each position there are two choices available to the model: go and no-go. Go is always chosen at the entrance to the Y-maze as bees that refuse to enter the maze would not continue the experiment. Following this there is a random choice of one of the two maze arms, left or right. If the model chooses no-go this procedure is repeated until the model chooses to go. As no learning

occurs at this stage it is possible for the model to constantly move between the two arms, never choosing to go. To prevent this eventuality we introduce a uniformly distributed random bias B to the go channel that increases with the number of times the model chooses no-go (N):

$$B = 10N(\mathcal{U}[0, 1] - 0.5).$$

The IN neurons are divided into non-overlapping groups of 8 neurons, each representing a stimulus. These are:

- Z: Stimulus for pretraining
- A: Stimulus for training pair
- B: Stimulus for training pair
- C: Stimulus for transfer test pair
- D: Stimulus for transfer test pair
- E: Stimulus for second transfer test pair
- F: Stimulus for second transfer test pair

Each group contains neurons which are zero when the stimulus is not present, and a value of $1 - \mathcal{U}[-0.05, 0.05]$ —consistent across the experiment for each bee, but not between bees—when active.

DMTS / DNMTS experimental procedure

Models as animals. We use the ‘models as animals’ principle to reproduce the experimental findings of [8], creating many variants of the model which perform differently in the experiments. To do this we change the random seed used to generate the connectivity c_{ij} between the IN and the KC neurons. For these experiments we use 360 model bee variants, each of which is individually tested, as this matches the number of bees in [8].

Pretraining familiarisation. As is undertaken in the experiments with real bees, we first familiarise our naive model bees with the experimental apparatus. This is done by first training ten rewarded repetitions of the bee entering the Y-maze with a stimulus not used in the experiment. In these cases the model does not choose between go and no-go, it is assumed that the first repetition represents the model finding the Y-maze and being heavily enough rewarded to complete the remaining repetitions. Following these ten repetitions the bee is trained with ten repetitions to travel to each of the two arms of the Y-maze. This procedure ensures that the bees will enter the maze and the two arms when the training begins, allowing them to learn the task.

Training. The training procedure comprises 60 trials in total, divided into blocks of 10 trials. The protocol involves a repeated set of four trials: two trials with each stimulus at the maze entrance, with each of these two trials having the stimulus at the maze entrance on different arms of the apparatus. In the case of match-to-sample the entrance stimulus is rewarded and the non-entrance stimulus is punished, and vice versa for not-match-to-stimulus.

Transfer test. For the transfer test we do not provide a reward or punishment, and test the models using the procedure for Training, substituting the transfer test stimuli for the training stimuli. Two sets of transfer stimuli are used, and four repetitions (left and right arm with each stimulus) are used for each set of stimuli.

Testing performance of the full model in other conditioning tasks

In addition to solving the DMTS and DNMTS tasks, we must validate that our proposed model can also perform a set of conditioning tasks that are associated with the mushroom

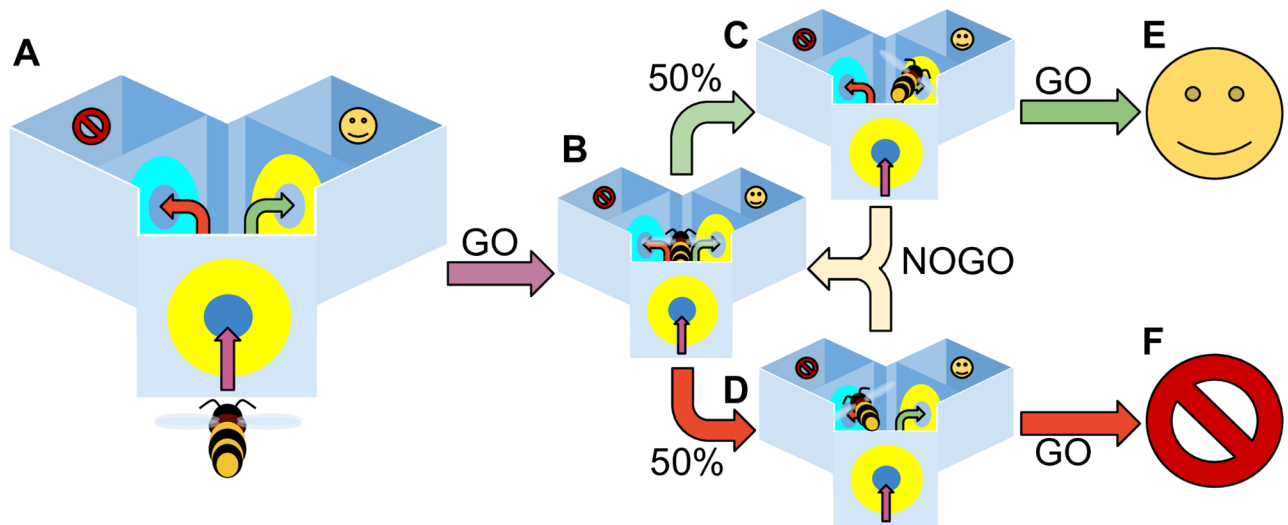


Fig 4. Experimental protocol for the model. The model bee is moved between a set of states which describe different locations in the Y-maze apparatus (A): at the entrance (B), in the central chamber facing the left arm (D), in the central chamber facing the right arm (D), in the left arm or in the right arm (E,F). When at the entrance or in the main chamber the bee is presented with a sensory input corresponding to one of the test stimuli; GO selection leads the bee to enter the maze when at the entrance, and to enter an arm and experience a potential reward when facing that arm; NOGO leads the bee to delay entering the maze, or to choose another maze arm uniformly at random, respectively. We can then set the test stimuli presented to match the requirements of a given trial (e.g. entrance (A), main chamber left (A), main chamber right (B) for DMTS when rewarding the left arm, or DNMTS when rewarding the right arm).

<https://doi.org/10.1371/journal.pcbi.1006435.g004>

bodies in bees, without our additional PCT circuits affecting performance. Importantly, these tasks are all performed with exactly the same model parameters that are used in the DMTS and DNMTS tasks, yet match the timescales and relative performances found in experiments performed on real bees. We choose four tasks, which comprise olfactory learning experiments using the proboscis extension reflex (PER) that are performed on restrained bees as well as visual learning experiments performed with free flying bees (Fig 4).

Differential learning / reversal experimental procedure (Fig 4, panel B). These experiments follow the same protocol as the DMTS experiments, except that for the first fifteen trials one stimulus is always rewarded when the associated arm is chosen (no reward or punishment is given for choosing the non-associated arm), and subsequent to trial fifteen the other stimulus is rewarded when the associated arm is chosen. No pretraining or transfer trials are performed and the data is analysed for each trial rather than in blocks of 10 due to the speed of learning acquisition. 200 virtual bees are used for this experiment (see Fig 4, panel B for results, to be compared with [47]).

Proboscis Extension Reflex (PER) experiments. The Proboscis Extension Reflex (PER) is a classical conditioning experimental paradigm used with restrained bees. In this paradigm the bees are immobilised in small metal tubes with only the head and antennae exposed. Olfactory stimuli (conditioned stimuli) are then presented to the restrained bees in association with a sucrose solution reward (unconditioned stimulus) (see [46] for full details).

For the PER experiments we separate the IN neurons as described in Section 1, however as the bees are restrained for these experiments we present odors following a pre-defined protocol, and the choices of the bee do not affect this protocol.

Single odor learning experimental procedure (Fig 4, panel A). In the single odor experiments we use the procedure outlined in Bitterman et al [46]. In this procedure acquisition and testing occur simultaneously. The real bees are presented an odor, and after a delay rewarded

with sucrose solution. If the animal extends its proboscis within the delay period it is rewarded directly and considered to have responded, if it does not the PER is invoked by touching the sucrose solution to the antennae and the animal is rewarded but considered not to have responded. To match this protocol the performance of the model was recorded at each trial, with NOGO considered a failure to respond to the stimulus, and GO a response. At each trial a reward was given regardless of the model's performance.

Positive / negative patterning learning experimental procedure (Fig 4, panels C & D). In these experiments we follow the protocol described in [48]. We divide the training into blocks, each containing four presentations of an odor or odor combination. For positive patterning we do not reward individual odors A and B, but reward the combination AB (A-B-AB+). In negative patterning we reward the odors A and B, but not the combination AB (A+B+AB-). In both cases the combined odor is presented twice for each presentation of the individual odors, so a block for positive patterning is [A-,AB+,B-,AB+] for example, while for negative patterning a block is [A+,AB-,B+,AB-]. Performance is assessed as for the single odor learning experiment, with the two combined odor responses averaged within each block.

Software and implementation

The reduced model was simulated in GNU Octave [64]. The full model was created and simulated using the SpineML toolchain [65] and the SpineCreator graphical user interface [66]. These tools are open source and installation and usage information can be found on the SpineML website at <http://spineml.github.io/>. Input vectors for the IN neurons and the state engine for navigation of the Y-maze apparatus are simulated using a custom script written in the Python programming language (Python Software Foundation, <https://www.python.org/>) interfaced to the model over a TCP/IP connection.

Statistical tests were performed as in [8] using 2×2 X^2 tests performed in R [67] using the `chisq.test()` function.

The code is available online at <http://github.com/BrainsOnBoard/bee-concept-learning>.

Acknowledgments

We thank Martin Giurfa, Thomas Nowotny and James Bennett for their constructive comments on the manuscript.

Author Contributions

Conceptualization: Alex J. Cope, Eleni Vasilaki.

Data curation: Alex J. Cope.

Formal analysis: Alex J. Cope, Eleni Vasilaki.

Funding acquisition: James A. R. Marshall, Andrew B. Barron.

Investigation: Alex J. Cope, Eleni Vasilaki, Dorian Minors.

Methodology: Alex J. Cope.

Project administration: James A. R. Marshall.

Software: Alex J. Cope.

Supervision: James A. R. Marshall.

Validation: Alex J. Cope, Eleni Vasilaki.

Visualization: Alex J. Cope.

Writing – original draft: Alex J. Cope, Eleni Vasilaki, James A. R. Marshall, Andrew B. Barron.

Writing – review & editing: Alex J. Cope, Eleni Vasilaki, Chelsea Sabo, James A. R. Marshall, Andrew B. Barron.

References

1. Wright AA, Katz JS. Generalization hypothesis of abstract-concept learning: Learning strategies and related issues in *Macaca mulatta*, *Cebus apella*, and *Columba livia*. *Journal of Comparative Psychology*. 2007; 121(4):387–397. <https://doi.org/10.1037/0735-7036.121.4.387> PMID: 18085922
2. Piaget J, Inhelder B. *The psychology of the child*. Basic books; 1969.
3. Daehler MW, Greco C. Memory in very young children. In: *Cognitive Learning and Memory in Children*. Springer New York; 1985. p. 49–79. Available from: http://link.springer.com/10.1007/978-1-4613-9544-7_2.
4. Avarguès-Weber A, Giurfa M. Conceptual learning by miniature brains. *Proceedings of the Royal Society of London B: Biological Sciences*. 2013; 280 (1772).
5. James W. *The principles of psychology*. vol. 1. Holt; 1890.
6. Wright AA. Concept learning and learning strategies. *Psychological Science*. 1997; 8(2):119–123. <https://doi.org/10.1111/j.1467-9280.1997.tb00693.x>
7. Wright AA. Testing the cognitive capacities of animals. *Learning and memory: The behavioral and biological substrates*. 1992; p. 45–60.
8. Giurfa M, Zhang S, Jenett A, Menzel R, Srinivasan MV. The concepts of 'sameness' and 'difference' in an insect. *Nature*. 2001; 410(6831):930–3. <https://doi.org/10.1038/35073582> PMID: 11309617
9. D'Amato MR, Salmon DP, Colombo M. Extent and limits of the matching concept in monkeys (*Cebus apella*). *Journal of Experimental Psychology: Animal Behavior Processes*. 1985; 11(1):35–51. PMID: 3989475
10. Diekamp B, Kalt T, Güntürkün O. Working memory neurons in pigeons. *The Journal of Neuroscience*. 2002;. <https://doi.org/10.1523/JNEUROSCI.22-04-j0002.2002>
11. Wallis JD, Anderson KC, Miller EK. Single neurons in prefrontal cortex encode abstract rules. *Nature*. 2001; 411(6840):953–956. <https://doi.org/10.1038/35082081> PMID: 11418860
12. Miller EK, Nieder A, Freedman DJ, Wallis JD. Neural correlates of categories and concepts. *Current opinion in neurobiology*. 2003; 13(2):198–203. [https://doi.org/10.1016/S0959-4388\(03\)00037-0](https://doi.org/10.1016/S0959-4388(03)00037-0) PMID: 12744974
13. Tomita H, Ohbayashi M, Nakahara K, Hasegawa I, Miyashita Y. Top-down signal from prefrontal cortex in executive control of memory retrieval. *Nature*. 1999; 401(6754):699–703. <https://doi.org/10.1038/44372> PMID: 10537108
14. Katz JS, Wright AA, Bodily KD. Issues in the comparative cognition of abstract-concept learning. *Comparative cognition & behavior reviews*. 2007; 2:79–92.
15. Baddeley AD, Hitch G. Working Memory. In: *Psychology of Learning and Motivation*. Elsevier; 1974. p. 47–89. Available from: <http://linkinghub.elsevier.com/retrieve/pii/S0079742108604521>.
16. Matsumoto Y, Menzel R, Sandoz JC, Giurfa M. Revisiting olfactory classical conditioning of the proboscis extension response in honey bees: A step toward standardized procedures. *Journal of Neuroscience Methods*. 2012; 211(1):159–167. <https://doi.org/10.1016/j.jneumeth.2012.08.018> PMID: 22960052
17. Arena P, Patané L, Stornanti V, Termini PS, Zäpf B, Strauss R. Modeling the insect mushroom bodies: Application to a delayed match-to-sample task. *Neural Networks*. 2013; 41:202–211. <https://doi.org/10.1016/j.neunet.2012.11.013> PMID: 23246431
18. Strausfeld NJ. *Arthropod brains: evolution, functional elegance, and historical significance*. Belknap Press of Harvard University Press Cambridge; 2012.
19. Menzel R. Searching for the memory trace in a mini-brain, the honeybee. *Learning & memory (Cold Spring Harbor, NY)*. 2001; 8(2):53–62. <https://doi.org/10.1101/lm.38801>
20. Søvik E, Perry CJ, Barron AB. Chapter Six—insect reward systems: comparing flies and bees. In: *Advances in Insect Physiology*. vol. 48; 2015. p. 189–226.

21. Giurfa M. Behavioral and neural analysis of associative learning in the honeybee: a taste from the magic well. *Journal of comparative physiology A, Neuroethology, sensory, neural, and behavioral physiology*. 2007; 193(8):801–24. <https://doi.org/10.1007/s00359-007-0235-9> PMID: 17639413
22. Galizia CG. Olfactory coding in the insect brain: data and conjectures. *European Journal of Neuroscience*. 2014; 39(11):1784–1795. <https://doi.org/10.1111/ejn.12558> PMID: 24698302
23. Mobbs PG. The brain of the honeybee *Apis Mellifera*. I. The connections and spatial organization of the mushroom bodies. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*. 1982; 298 (1091). <https://doi.org/10.1098/rstb.1982.0086>
24. Boitard C, Devaud JM, Isabel G, Giurfa M. GABAergic feedback signaling into the calyces of the mushroom bodies enables olfactory reversal learning in honey bees. *Frontiers in Behavioral Neuroscience*. 2015; 9. <https://doi.org/10.3389/fnbeh.2015.00198> PMID: 26283938
25. Devaud JM, Papouin T, Carcaud J, Sandoz JC, Grünwald B, Giurfa M. Neural substrate for higher-order learning in an insect: mushroom bodies are necessary for configural discriminations. *Proceedings of the National Academy of Sciences of the United States of America*. 2015; 112(43):E5854–62. <https://doi.org/10.1073/pnas.1508422112> PMID: 26460021
26. Bazhenov M, Huerta R, Smith BH. A computational framework for understanding decision making through integration of basic learning rules. *Journal of Neuroscience*. 2013; 33(13). <https://doi.org/10.1523/JNEUROSCI.4145-12.2013>
27. Huerta R, Nowotny T. Fast and robust learning by reinforcement signals: explorations in the insect brain. *Neural Computation*. 2009; 21(8):2123–2151. <https://doi.org/10.1162/neco.2009.03-08-733> PMID: 19538091
28. Fahrbach SE. Structure of the mushroom bodies of the insect brain. *Annual Review of Entomology*. 2006; 51(1):209–232. <https://doi.org/10.1146/annurev.ento.51.110104.150954> PMID: 16332210
29. Heisenberg M. Mushroom body memoir: from maps to models. *Nature Reviews Neuroscience*. 2003; 4(4):266–275. <https://doi.org/10.1038/nrn1074> PMID: 12671643
30. Schwaerzel M, Monastirioti M, Scholz H, Friggi-Grelin F, Birman S, Heisenberg M. Dopamine and octopamine differentiate between aversive and appetitive olfactory memories in *Drosophila*. *Journal of Neuroscience*. 2003; 23(33). <https://doi.org/10.1523/JNEUROSCI.23-33-10495.2003> PMID: 14627633
31. Strube-Bloss MF, Nawrot MP, Menzel R. Mushroom body output neurons encode odor-reward associations. *Journal of Neuroscience*. 2011; 31(8). <https://doi.org/10.1523/JNEUROSCI.2583-10.2011> PMID: 21414933
32. Szyszka P, Galkin A, Menzel R. Associative and non-associative plasticity in kenyon cells of the honeybee mushroom body. *Frontiers in systems neuroscience*. 2008; 2:3. <https://doi.org/10.3389/neuro.06.003.2008> PMID: 18958247
33. Hattori D, Aso Y, Swartz KJ, Rubin GM, Abbott LF, Axel R. Representations of Novelty and Familiarity in a Mushroom Body Compartment. *Cell*. 2017; 169(5):956–969.e17. <https://doi.org/10.1016/j.cell.2017.04.028> PMID: 28502772
34. Tsodyks MV, Markram H. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*. 1997; 94(2):719–723. <https://doi.org/10.1073/pnas.94.2.719>
35. Vasilaki E, Giugliano M. Emergence of connectivity motifs in networks of model neurons with short-and long-term plastic synapses. *PLoS One*. 2014; 9(1):e84626. <https://doi.org/10.1371/journal.pone.0084626> PMID: 24454735
36. Esposito U, Giugliano M, Vasilaki E. Adaptation of short-term plasticity parameters via error-driven learning may explain the correlation between activity-dependent synaptic properties, connectivity motifs and target specificity. *Frontiers in computational neuroscience*. 2014; 8. <https://doi.org/10.3389/fncom.2014.00175> PMID: 25688203
37. Yarden TS, Nelken I. Stimulus-specific adaptation in a recurrent network model of primary auditory cortex. *PLoS computational biology*. 2017; 13(3):e1005437. <https://doi.org/10.1371/journal.pcbi.1005437> PMID: 28288158
38. Ganeshina O, Menzel R. GABA-immunoreactive neurons in the mushroom bodies of the honeybee: an electron microscopic study. *The Journal of comparative neurology*. 2001; 437(3):335–49. <https://doi.org/10.1002/cne.1287> PMID: 11494260
39. Haehnel M, Menzel R. Sensory representation and learning-related plasticity in mushroom body extrinsic feedback neurons of the protocerebral tract. *Frontiers in Systems Neuroscience*. 2010; 4:161. <https://doi.org/10.3389/fnsys.2010.00161> PMID: 21212833
40. Rybak J, Menzel R. Anatomy of the mushroom bodies in the honey bee brain: the neuronal connections of the alpha-lobe. *The Journal of comparative neurology*. 1993; 334(3):444–65. <https://doi.org/10.1002/cne.903340309> PMID: 8376627

41. Okada R, Rybak J, Manz G, Menzel R. Learning-related plasticity in PE1 and other mushroom body-extrinsic neurons in the honeybee brain. *Journal of Neuroscience*. 2007; 27(43). <https://doi.org/10.1523/JNEUROSCI.2216-07.2007>
42. Perry CJ, Barron AB. Neural mechanisms of reward in insects. *Annual Review of Entomology*. 2013; 58(1):543–562. <https://doi.org/10.1146/annurev-ento-120811-153631> PMID: 23020615
43. Papadopoulou M, Cassenaer S, Nowotny T, Laurent G. Normalization for Sparse Encoding of Odors by a Wide-Field Interneuron. *Science*. 2011; 332 (6030). <https://doi.org/10.1126/science.1201835> PMID: 21551062
44. Huerta R, Nowotny T, García-Sánchez M, Abarbanel HDI, Rabinovich MI. Learning Classification in the Olfactory System of Insects. *Neural Computation*. 2004; 16(8):1601–1640. <https://doi.org/10.1162/089976604774201613> PMID: 15228747
45. Peng F, Chittka L. A simple computational model of the bee mushroom body can explain seemingly complex forms of olfactory learning and memory. *Current biology: CB*. 2016; 0(0):2597–2604.
46. Bitterman ME, Menzel R, Fietz A, Schäfer S. Classical conditioning of proboscis extension in honeybees (*Apis mellifera*). *Journal of comparative psychology (Washington, DC: 1983)*. 1983; 97(2):107–19.
47. Giurfa M. Conditioning procedure and color discrimination in the honeybee *Apis mellifera*. *Naturwissenschaften*. 2004; 91(5):228–231. <https://doi.org/10.1007/s00114-004-0530-z> PMID: 15146270
48. Deisig N, Lachnit H, Giurfa M, Hellstern F. Configural olfactory learning in honeybees: negative and positive patterning discrimination. *Learning & memory (Cold Spring Harbor, NY)*. 2001; 8(2):70–8. <https://doi.org/10.1101/lm.8.2.70>
49. Moore TL, Schettler SP, Killiany RJ, Rosene DL, Moss MB. Impairment in delayed nonmatching to sample following lesions of dorsal prefrontal cortex. *Behavioral neuroscience*. 2012; 126(6):772–80. <https://doi.org/10.1037/a0030493> PMID: 23088539
50. Avarguès-Weber A, D’Amaro D, Metzler M, Dyer AG. Conceptualization of relative size by honeybees. *Frontiers in Behavioral Neuroscience*. 2014; 8:80. <https://doi.org/10.3389/fnbeh.2014.00080> PMID: 24672444
51. Zhang S, Bock F, Si A, Tautz J, Srinivasan MV. Visual working memory in decision making by honey bees. *Proceedings of the National Academy of Sciences of the United States of America*. 2005; 102(14):5250–5. <https://doi.org/10.1073/pnas.0501440102> PMID: 15795382
52. Lind J, Enquist M, Ghirlanda S. Animal memory: A review of delayed matching-to-sample data. *Behavioural Processes*. 2015; 117:52–58. <https://doi.org/10.1016/j.beproc.2014.11.019> PMID: 25498598
53. Wright AA, Katz JS. Mechanisms of same/different concept learning in primates and avians. *Behavioural Processes*. 2006; 72(3):234–254. <https://doi.org/10.1016/j.beproc.2006.03.009> PMID: 16621333
54. Seger CA, Miller EK. Category learning in the brain. *Annual Review of Neuroscience*. 2010; 33(1):203–219. <https://doi.org/10.1146/annurev-neuro.051508.135546> PMID: 20572771
55. Tsujimoto S, Genovesio A, Wise SP. Comparison of strategy signals in the dorsolateral and orbital prefrontal cortex. *Journal of Neuroscience*. 2011; 31(12):4583–4592. <https://doi.org/10.1523/JNEUROSCI.5816-10.2011> PMID: 21430158
56. Premack D. On the abstractness of human concepts: Why it would be difficult to talk to a pigeon. *Cognitive processes in animal behavior*. 1978; p. 423–451.
57. Premack D, Premack AJ. *The mind of an ape*; 1983.
58. Thompson RK, Oden DL. A profound disparity revisited: Perception and judgment of abstract identity relations by chimpanzees, human infants, and monkeys. *Behavioural processes*. 1995; 35(1-3):149–61. [https://doi.org/10.1016/0376-6357\(95\)00048-8](https://doi.org/10.1016/0376-6357(95)00048-8) PMID: 24896027
59. Lloyd Morgan C. *An introduction to comparative psychology*. London: W Scott Publishing Co; 1903.
60. Sober E. *Ockham’s razors*. Cambridge University Press; 2015.
61. Mikhalevich I. Experiment and Animal Minds: Why the Choice of the Null Hypothesis Matters. *Philosophy of Science*. 2015; 82(5):1059–1069. <https://doi.org/10.1086/683440>
62. Vasilaki E, Frémaux N, Urbanczik R, Senn W, Gerstner W. Spike-Based Reinforcement Learning in Continuous State and Action Space: When Policy Gradient Methods Fail. *PLoS Computational Biology*. 2009; 5(12):e1000586. <https://doi.org/10.1371/journal.pcbi.1000586> PMID: 19997492
63. Richmond P, Buesing L, Giugliano M, Vasilaki E. Democratic Population Decisions Result in Robust Policy-Gradient Learning: A Parametric Study with GPU Simulations. *PLoS ONE*. 2011; 6(5):e18539. <https://doi.org/10.1371/journal.pone.0018539> PMID: 21572529
64. John W Eaton David Bateman SH, Wehbring R. {GNU Octave} version 4.0.0 manual: a high-level interactive language for numerical computations. GNU Octave; 2015. Available from: <http://www.gnu.org/software/octave/doc/interpreter>.

65. Richmond P, Cope A, Gurney K, Allerton DJ. From model specification to simulation of biologically constrained networks of spiking neurons. *Neuroinformatics*. 2013; 12(2):307–23. <https://doi.org/10.1007/s12021-013-9208-z>
66. Cope AJ, Richmond P, James SS, Gurney K, Allerton DJ. SpineCreator: A graphical user interface for the creation of layered neural models. In-press. 2015
67. R Core Team. R: A Language and Environment for Statistical Computing; 2013. Available from: <http://www.r-project.org/>.