



# Putting a bug in ML: The moth olfactory network learns to read MNIST

Charles B. Delahunt<sup>a,b,\*</sup>, J. Nathan Kutz<sup>a</sup>

<sup>a</sup> Department of Applied Mathematics, University of Washington, Seattle, United States

<sup>b</sup> Computational Neuroscience Center, University of Washington, Seattle, United States

## ARTICLE INFO

### Article history:

Received 5 October 2018

Received in revised form 15 May 2019

Accepted 19 May 2019

Available online 4 June 2019

### Keywords:

Hebbian

Sparsity

Olfactory network

Bio-mimesis

Neural networks

One-shot learning

## ABSTRACT

We seek to (i) characterize the learning architectures exploited in biological neural networks for training on very few samples, and (ii) port these algorithmic structures to a machine learning context. The moth olfactory network is among the simplest biological neural systems that can learn, and its architecture includes key structural elements and mechanisms widespread in biological neural nets, such as cascaded networks, competitive inhibition, high intrinsic noise, sparsity, reward mechanisms, and Hebbian plasticity. These structural biological elements, in combination, enable rapid learning.

MothNet is a computational model of the moth olfactory network, closely aligned with the moth's known biophysics and with *in vivo* electrode data collected from moths learning new odors. We assign this model the task of learning to read the MNIST digits. We show that MothNet successfully learns to read given very few training samples (1–10 samples per class). In this few-samples regime, it outperforms standard machine learning methods such as nearest-neighbors, support-vector machines, and neural networks (NNs), and matches specialized one-shot transfer-learning methods but without the need for pre-training.

The MothNet architecture illustrates how algorithmic structures derived from biological brains can be used to build alternative NNs that may avoid the high training data demands of many current engineered NNs.

© 2019 Elsevier Ltd. All rights reserved.

## 1. Introduction

Originally inspired by the biological structure of networks of interacting neurons (Fukushima, 1980; Hubel & Wiesel, 1962; McCulloch & Pitts, 1943; Rosenblatt, 1958), neural networks (NNs) have since developed a suite of algorithmic tools (such as backprop, convolutional kernels, etc.) which, combined into complex and deep NN architectures, have achieved unprecedented success in a wide array of machine learning (ML) tasks (Goodfellow, Bengio, & Courville, 2016; LeCun, Bengio, & Hinton, 2015; Schmidhuber, 2015). However, they have trouble learning from few samples. We seek to improve NN performance on such tasks by revisiting the well of biological example to characterize key biological structures involved in learning, for transfer to the NN context.

In this work, we apply MothNet, a NN architecture closely based on the *Moth Olfactory Network*, to the task of learning vectorized MNIST digits. MothNet successfully learns to read given very few training samples (1–10 samples per class) and in fact outperforms standard machine learning methods such as

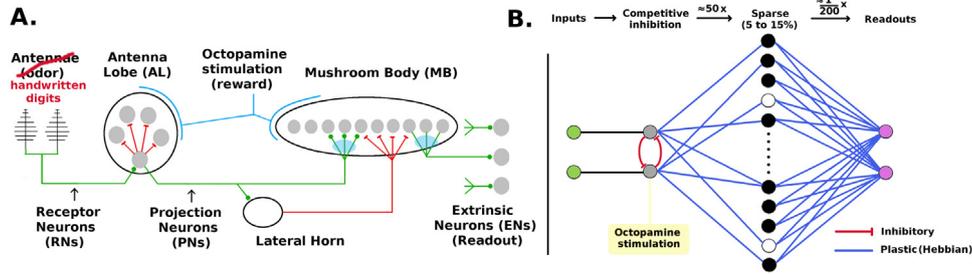
Nearest-Neighbors, support-vector machines (SVMs), Neural Nets (i.e. standard fully-connected, feed-forward, non-convolutional nets with softmax loss functions), and convolutional neural networks (CNNs), in the few-samples regime (Fig. 5). In addition, it performs comparably to specialized one-shot learning NNs, but without their need for pre-training to initialize NN weights (Table 1).

The moth olfactory network is among the simplest biological neural networks (BNNs) that can learn (Riffell, Lei, Abrell, & Hildebrand, 2012).<sup>1</sup> It is well-characterized, and it contains key features widespread in BNNs, including high noise (Galizia, 2014), random connections (Caron, Ruta, Abbott, & Axel, 2013), competitive inhibition (Hong & Wilson, 2015), Hebbian synaptic growth (Cassenaer & Laurent, 2007; Masse, Turner, & Jefferis, 2009), high-dimensional sparse layers (Campbell & Turner, 2010; Honegger, Campbell, & Turner, 2011), large dimension shifts between layers (Babadi & Sompolinsky, 2014), and generalized stimulation of neurons during learning (Hammer & Menzel, 1998). It thus

<sup>1</sup> References here include studies of various insect species. MothNet is based on the moth olfactory network in particular, but draws on studies from other insects to fill gaps where studies on the moth are lacking, subject to the constraint that inter-species transferability varies (the fly and moth are similar, while locust and honeybee results are sometimes general enough to transfer, sometimes not).

\* Corresponding author at: Department of Applied Mathematics, University of Washington, Seattle, United States.

E-mail addresses: [delahunt@uw.edu](mailto:delahunt@uw.edu) (C.B. Delahunt), [kutz@uw.edu](mailto:kutz@uw.edu) (J.N. Kutz).



**Fig. 1.** Network schematics. **A:** Biological schematic. Green lines show excitatory connections, red lines show inhibitory connections. Light blue ovals show plastic connections into and out of the MB. The units in the AL competitively inhibit each other. Global inhibition from the lateral horn induces sparsity on MB responses. The ENs give the final, actionable readouts of the system's response to a stimulus. **B:** Alternative schematic of the same network. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

offers an ideal avenue to investigate biological learning and the interacting components that make learning possible.

The olfactory processing unit centers on two interacting networks, the noisy antennal lobe (AL) and the sparse mushroom body (MB), known as the AL–MB (Campbell & Turner, 2010; Wilson, 2008). The AL pre-processes odor inputs, then projects the odor signal forward to the MB, which in turn feeds forward to readout neurons. Learning occurs when the neuromodulatory chemical octopamine (triggered by sucrose reward) induces an overall increase in excitation in the AL (Dacks, Riffell, Martin, Gage, & Nighorn, 2012), causing the plastic MB synapses to update via Hebbian (“fire together, wire together”) growth.

A computational model of a stand-alone honeybee MB (i.e. without the AL or octopamine dynamics) was trained by Huerta and Nowotny (2009) on the handwritten vectorized MNIST digits dataset (LeCun & Cortes, 2010). The trained honeybee MB attained over 80% accuracy given sufficient training samples (200–2000 samples per class), indicating that the MB is to some degree task-agnostic, i.e. not limited solely to odor processing.

MothNet is an end-to-end computational model of the *Manduca sexta* (Hawk moth) olfactory network developed by Delahunt, Riffell, and Kutz (2018b) to study the moth's olfactory learning mechanisms. MothNet's architecture is closely based on the moth's known biophysical structure. It includes the AL, the MB, and readout neurons, as well as Hebbian plasticity and neuromodulatory stimulation by octopamine. Its behavior is calibrated to, and consistent with, *in vivo* firing rate data recorded in moths learning new odors. For full details see Delahunt et al. (2018b).

To test the generalizability of this learning architecture, we assigned MothNet the ML task of identifying downsampled, vectorized MNIST digits (hereafter “*v*MNIST” to emphasize that the inputs lack spatial structure). We made minimal modifications to the MothNet architecture: We simply replaced odor inputs to the AL neurons with pixel values, increased the number of AL units from 60 to 85, and attached 10 readout neurons, one for each *v*MNIST digit. MothNet routinely achieved 70%–80% accuracy classifying test digits after training on 1–10 samples per class, out-performing standard ML methods. These results demonstrate that even a very simple biological architecture contains novel and effective tools that are useful for ML tasks, in particular tasks constrained by limited training data or the need to add and train new classes without retraining the full NN.

## 2. The moth olfactory network and MothNet model

### 2.1. Moth olfactory network outline

The network is organized as a feed-forward cascade of five distinct networks, as well as a reward mechanism (Kvellido, Løfaldli, Rybak, Menzel, & Mustaparta, 2009; Martin et al., 2011). Fig. 1 gives a system schematic.

Starting at the Antennae, several thousand noisy chemical Receptor Neurons (RNs) detect odor and send signals to the Antennal Lobe (AL) (Masse et al., 2009; Wilson, 2008). Most of these RNs are narrowly tuned (to 1 or 2 chemicals), while others are more broadly tuned (Shields & Hildebrand, 2001). Each RN type converges onto one of roughly 60 isolated units (glomeruli) in the AL (Martin et al., 2011). The AL acts as a pre-amp, providing gain control and sharpening odor codes through lateral inhibition (Bhandawat, Olsen, Gouwens, Schlieff, & Wilson, 2007). The AL projects odor codes forward to the Mushroom Body (MB) (Campbell & Turner, 2010), in excitatory Projection Neurons (PNs) that randomly connect (non-densely) to MB neurons (Caron et al., 2013).

The MB contains about 4000 Kenyon Cells, which fire sparsely and encode odor signatures as memories (Honegger et al., 2011; Perisse, Burke, Huetteroth, & Waddell, 2013) (we use “MB” as a synonym for these Kenyon Cells). Sparsity in the MB is enforced by global inhibition from the Lateral Horn (Bazhenov & Stopfer, 2010). The MB feeds forward to Extrinsic Neurons (ENs), numbering  $\sim 10$ 's, which are “readout neurons” that interpret the MB codes and deliver actionable output to the rest of the moth (Campbell et al., 2013; Hige, Aso, Rubin, & Turner, 2015).

The network can learn: In response to reward (sugar at the proboscis), several neurons release the neuromodulator octopamine globally in the AL and MB (Dacks, Christensen, Agricola, Wollweber, & Hildebrand, 2005) (for dopamine (Dacks et al., 2012); for honeybees (Hammer & Menzel, 1995, 1998)). Octopamine stimulates AL neurons, inciting higher firing rate responses to the reinforced stimuli (e.g. the *in vivo* data used in Delahunt et al., 2018b and Dacks et al., 2012 for dopamine). This in turn induces growth in the plastic synaptic connections into the MB (AL→MB) and out of the MB (MB→ENs) via Hebbian updates (Cassenaer & Laurent, 2007; Masse et al., 2009). Learning does not occur without this octopamine input (Hammer & Menzel, 1995, 1998).

### 2.2. MothNet model

The MothNet computational model closely follows this biological architecture in terms of connections, numbers of neurons in each layer, etc. (Delahunt et al., 2018b). Neural firing rates are modeled with firing-rate dynamics (Dayan & Abbott, 2005) evolved as stochastic differential equations (SDEs) (Higham, 2001):

$$\tau \frac{dx}{dt} = -x + s(\Sigma \mathbf{w}_i \mathbf{u}_i) = -x + S(\mathbf{w} \cdot \mathbf{u}) + dW, \quad (1)$$

where  $x(t)$  = firing rate (FR) for a neuron;  $\mathbf{w}$  = connection weights;  $\mathbf{u}$  = upstream neuron FRs;  $S(\cdot)$  is a sigmoid function or similar; and  $W(t)$  = a Brownian motion process.

In MothNet this equation is modified by an inserted term to model the stimulative effect of octopamine on the AL during learning, since this effect is central to learning in the computational model (Delahunt et al., 2018b). The model uses a simple model of Hebbian plasticity for synaptic weight updates (Dayan & Abbott, 2005; Hebb, 1949; Roelfsema & Holtmaat, 2018):

$$\Delta w_{ab}(t) = \gamma f_a(t) f_b(t) \quad (2)$$

where  $f_a(t), f_b(t)$  are firing rates of neurons  $a, b$  at time  $t$ ;  $w_{ab}$  is the synaptic weight between them; and  $\gamma$  is a growth rate parameter.

In addition, inactive MB→EN weights are subject to proportional decay:

$$\Delta w_{ab}(t) = -\delta w_{ab}(t), \text{ if } f_a(t) f_b(t) = 0. \quad (3)$$

where  $\delta$  is a decay parameter. There are two layers of plastic synaptic weights: AL→MB, and MB→ENs (i.e. pre- and post-MB). Hebbian plasticity is assumed to be “switched on” by reward, so it is cosynchronous with octopamine. Thus, plasticity only occurs during training sessions. There is no built-in stabilization of weight updates, except the upper and lower (i.e. 0) rails of allowable weights. As seen in Fig. 7, sparsity in the MB is crucial to prevent Hebbian weight updates from running out of control and amplifying noise channels.

For the  $v$ MNIST task, MothNet has 10 extrinsic neurons (ENs), that start with identical MB→EN connections. When training begins, these 10 ENs are randomly assigned to target particular digits 0 to 9. Once ENs are assigned, training is supervised: When a digit of class  $j$  is presented, only MB→EN $_j$  connections are updated, where EN $_j$  is the EN assigned to class  $j$ . That is, the system knows the class of the training sample, for the purposes of post-MB updates. Training rapidly individualizes these weights according to their target digits. In contrast, all pre-MB connections are updated in every case, since these connections are common to all inputs (in these experiments, pre-MB connections were essentially fixed due to slow learning rates).

We note that in MothNet, updates do not follow the same logic as in typical ML algorithms: There is no objective function to be minimized, and no output-based loss that is pushed back through the network as with backprop or agent-based reinforcement learning (there is no “agent” in the MothNet system). The Hebbian weight updates, either growth or decay, occur on a purely local “use it or lose it” basis.

### 3. Methods

#### 3.1. Training data

We did not use the MNIST dataset in its usual format of images with spatial structure. Rather, we derived an 85-feature, 10-class dataset, with no spatial information, from the MNIST dataset. MNIST images (from PMTK3, Murphy, 2012) were downsampled and vectorized, and pixel values then served as input features. Of course both biological and engineered systems routinely choose better feature sets. However, pixels-as-features provided a good test of whether MothNet can effectively learn to discriminate classes given inputs with inter-class correlations.

MothNet maps features to units in the AL in 1-to-1 fashion (a simplification of biological reality). Using full MNIST images, vectorized to give pixels-as-features, would imply  $28^2 = 784$  AL units. To keep the scale of MothNet somewhat close to that of the actual moth (60 AL units), we preprocessed the images as follows:

1. Crop by 2 and downsample by 2 (linear interpolation), to get  $12 \times 12$  images (shown in Fig. 2).

2. Mean-subtract using 500 random, set-aside digits (50 from each class) and zero out negative values.

3. Select the most-active pixels by thresholding the various class averages. The purpose is to retain the most globally active pixels while also preserving the most active pixels of each class, and to exclude border pixels which supply little information. This step was part of the dataset definition, not part of training, and did not involve the MothNet model.

4. Vectorize the remaining pixels.

Each retained pixel became a feature that fed into one unit of MothNet’s AL. The experiments described here used 85 pixels (out of  $12^2 = 144$  total) to represent the  $v$ MNIST digits. This gave MothNet 85 units in the AL and 2550 neurons in the MB (i.e. 30 times the number of AL units). Examples of downsampled images, before vectorization and restriction to the receptive field pixels, are shown in Fig. 2.

#### 3.2. Experiment design

Instances of MothNet are randomly generated from parameter templates. Particular MothNet behaviors can be modulated by varying one or two template parameters. Each experiment used a fixed basic template, modified only by varying the parameter(s)-under-test. Each of these modified templates then randomly generated many MothNet instances, typically 13–17 per data point. The basic templates varied slightly between experiments, e.g. in some parameter values. We found that these slight differences in template had minor effect, and that a wide range of parameters and templates delivered effective learning behavior. Each experiment contained three stages:

1. Pre-training Baseline (15 digits per class), used to assess naive classifier accuracy;
2. Training ( $N$  digits per class, randomly ordered);
3. Post-Training Validation (15 digits per class), used to assess post-training classifier accuracy.

Digits for each stage were randomly chosen without replacement from non-intersecting pools. Using more than 15 digits in baseline and validation sets did not significantly affect results.

Since the moth olfactory system can learn to recognize a new odor given roughly 8–10 samples (Riffell et al., 2012), we focused on small training sets (1–20 samples/class). In some experiments, training included multiple “sniffs”, i.e. repeated presentations of each sample. The order of training samples did not matter, perhaps because the strongest plasticity was specific to the readout neuron (EN) targeting a given class.

#### 3.3. Classifiers

System readout units are the Extrinsic Neurons (ENs) downstream from the sparse MB layer and its plastic connections. These ENs are silent absent any input sample, and they consistently respond, more or less strongly, to input samples (see Fig. 3). We use two different classifiers to assess MothNet performance, both dependent on the EN readouts.

##### 3.3.1. Softmax classifier

The first method is the basic softmax:

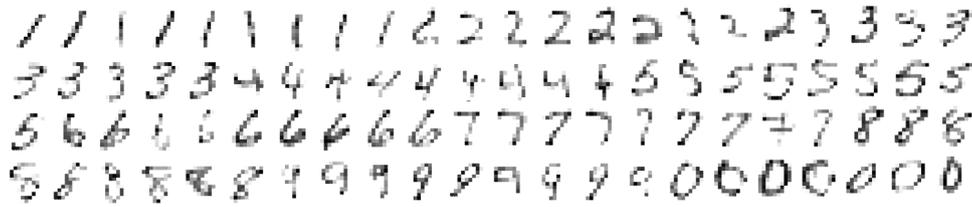
$$\hat{s} = \max_{j \in J} \left\{ \frac{e^{E_j(s)}}{\sum_{i \in J} e^{E_i(s)}} \right\}, \text{ where} \quad (4)$$

$\hat{s}$  = predicted class of sample  $s$

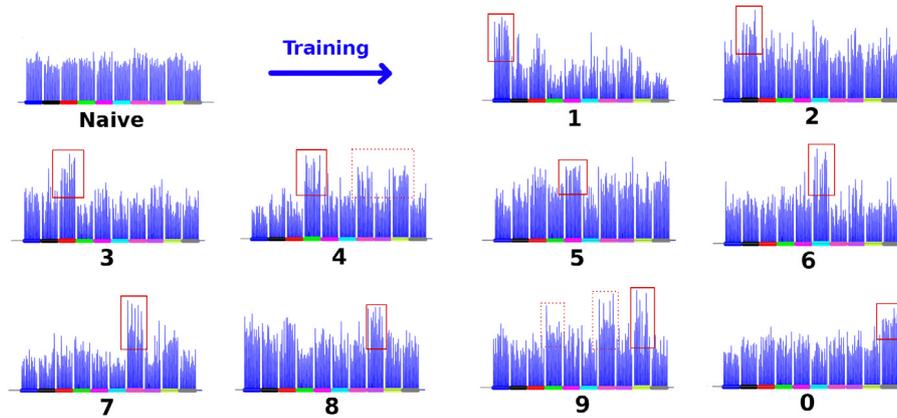
$E_i(s)$  = response of the  $i$ th EN to  $s$

$j \in J$  are the classes (0–9).

This is tantamount to prediction based on the strongest EN response. In NNs, the softmax rule is typically baked into the loss function, so it is a logical classifier to use, by construction.



**Fig. 2.** Cropped, downsampled, mean-subtracted and zero-thresholded MNIST digit images (pre-processing steps 1 and 2), shown before further sub-selection of globally most-active pixels and vectorization (steps 3 and 4). Vectors of these pixels were the feature vectors for the experiments.



**Fig. 3.** EN time courses for a typical MothNet experiment before and after training, showing how training causes the class response distributions to diverge in each EN. Each block of spikes shows the time-course of an EN's response to 150 digits (15 ones, then 15 twos, etc.). The variation in response (spike height) to digits of a single class reflect variations in the digits and also effects of AL noise. The top left block shows the naive responses (all ENs had similar responses). The blue arrow represents the training stage. The numbered blocks show the responses of the various ENs to 150 validation digits after training has completed. The EN's responses to digits in the targeted class are framed in red, and some confounding class responses in dashed red). This figure shows results of time-evolutions from one MothNet instance. The equivalent statistical results, over several MothNet instances, are shown in Fig. 4.

However, we see no reason that it would be well-suited to a network that uses Hebbian updates. In particular, a Hebbian update method does not constrain the scales of the various EN responses to match each other, as happens in a standard NN with softmax-based loss function.

### 3.3.2. Log-likelihood classifier

We thus use a second classifier as well, which tries to capture the combined statistical content of the various EN readouts without assuming that the EN responses match each other in scale. This classifier is a summed log-likelihood over the distributions of responses to each digit class in each EN:

$$\hat{s} = \min_{j \in J} \left\{ \sum_{i \in J} \left( \frac{E_i(s) - \mu E_{ij}}{\sigma E_{ij}} \right)^4 \right\}, \text{ where} \quad (5)$$

$\hat{s}$  = predicted class of sample  $s$

$E_i(s)$  = response of the  $i$ th EN to  $s$

$\mu E_{ij}$  =  $\text{mean}(E_i(t) | t \in V \text{ and } t \in \text{class } j)$

$\sigma E_{ij}$  =  $\text{std dev}(E_i(t) | t \in V \text{ and } t \in \text{class } j)$

$j \in J$  are the classes (0–9)

$V$  is a reference set (e.g. a validation set).

Roughly,  $j$  is a strong candidate for  $\hat{s}$  if each EN's response to  $s$  is close (in Mahalanobis distance) to that EN's expected response to class  $j$ . The use of the 4<sup>th</sup> power (vs the usual 2<sup>nd</sup> power) is a sharpener that penalizes outliers.

This summed log-likelihood is a measure of how well the ENs can separate the response distributions to various classes, combining information from all ENs and including information about responses to classes not specifically targeted by a particular EN. The goal of this classifier is to assess how much discriminatory

information MothNet was able to extract from the training data. We do not wish to imply it is biologically realistic (we do not know). Accuracy of naive (i.e. untrained) MothNet instances was about 15% using the log-likelihood classifier, slightly higher than random guessing, perhaps because the digit “1” often elicited slightly different naive responses than other digits.

For a given experiment, the post-training classification accuracy was calculated on the validation set, i.e. the same set was used to estimate the post-training EN response distribution parameters  $\mu E$  and  $\sigma E$ . Similarly, the baseline (pre-training) classification accuracy was calculated on the same baseline set used to estimate naive EN response distribution parameters. Holdout sets had similar accuracy.

### 3.3.3. Pros and cons

The softmax classifier does not require a validation set, and thus is a more realistic candidate for true few-shot learning, where by definition no validation set is available. However, it is an arbitrary choice for systems that use a Hebbian learning rule, especially since the Hebbian rule does not ensure a built-in means of scaling EN responses. Thus softmax may be ill-suited to reflect the full information of the trained network effectively.

The Log-likelihood classifier can arguably claim to better leverage the information contained in a Hebbian-trained network, and it is robust to differences in scale of EN responses. However, it requires a validation set for its hyperparameters. In few-shot scenarios, while it may well express the stored learning of the network it is not a practical in-the-field classifier.

We find that in practice, each classifier has strengths. In particular, the log-likelihood is stronger given few training samples, while softmax is stronger given many training samples. We report results for both.

## 4. Experiments and results

This section first presents (i) results of MothNet experiments focused on learning. It next gives baseline comparisons of MothNet to (ii) standard ML methods, and (iii) specialized ML one-shot methods. It then presents further results of MothNet experiments focused on (iv) one-shot learning, (v) sniffing, (vi) effects of AL noise, and (vii) MB sparsity. Computer code for running simulations and experiments in this paper can be found at: <https://github.com/charlesDelahunt/PuttingABugInML>

### 4.1. Learning experiments

We ran training experiments with various MothNet templates to assess their ability to learn the  $\nu$ MNIST digits. In general, a wide range of templates responded well to training by differentiating their EN responses to different digits (input classes). In naive (untrained) MothNet instances, all ENs had the same response profile, and responded similarly to all digit classes, as expected given the symmetry of random connection weights to the various ENs. Training caused EN responses to rapidly diverge from baseline and from each other, such that each EN responded most strongly to its assigned digit. Common effects of training included:

1. Most ENs (e.g. 1, 2, 6, 7, 0) tended to amplify the response to their trained digit very well, compared to responses to control digits. This gave strong separation and accurate classification.
2. A few ENs (e.g. 5) sometimes poorly separated their trained digit from control digits. These were the digits most often misclassified during validation.
3. Some ENs consistently boosted the responses to certain control digits along with their trained digit (e.g., EN<sub>9</sub> boosted 4 and 7, EN<sub>4</sub> boosted 7 and 9). These cases typically reflected visible similarities in the digits, and led to characteristic errors. For example, 9s, if misclassified, were often misclassified as 4s. However, 9s were not misclassified as 7s, because EN<sub>7</sub> usually strongly separated 7 from 9. That is, outputs of EN<sub>9</sub>, EN<sub>7</sub>, and EN<sub>4</sub> combined were sufficient to distinguish 9 from 7, but not always 9 from 4.

These behaviors are evident in Fig. 3, which shows time-courses of EN firing rate responses, pre- and post-training, for a typical MothNet instance. Each subplot shows EN response to 150 digits (15 ones, then 15 twos, etc.). The post-training responses are normalized by the EN's mean trained class response for clarity. Training typically increased and/or decreased all class responses of an EN, but to different degrees, resulting in the separations seen in the normalized timecourses. Mid-training responses are excised from the timecourses to save space; these responses were consistently much stronger due to the stimulating effect of octopamine injected during training, and would extend past the top of the plot.

Fig. 4 plots EN response distribution statistics (mean  $\pm$  std dev) from a typical experiment, in which 13 MothNet instances were generated from a template then trained on 15 samples per class. Post-training accuracy (log-likelihood) for MothNet instances of this template was 71%–83%, starting from  $\sim$ 15% baseline accuracy. Similar results, both accuracies and limitations (such as confusing 4s and 9s) held for a wide range of parameter templates and training regimes. These learning experiments indicate that a model closely based on the moth olfactory network can rapidly learn to read handwritten digits. This aptitude was qualified by an apparent upper limit of about 80% on mean population accuracy (though gifted individuals attained up to 85% accuracy), given such constraints as the restriction to “natural” model parameters, downsampled images, and the use of vectorized (non-spatial) input features.

### 4.2. Comparison to standard ML methods

To set a learning performance baseline, we trained four standard ML methods (Nearest Neighbors, SVM, Neural Net, and CNN) in the limited-training-data regime ( $N \leq 100$  samples per class), using the same downsampled, non-spatial dataset (except CNN, see below). Hyperparameter values can strongly affect algorithm performance. We applied due diligence in optimizing hyperparameters for these ML models, and we feel our optimizations are strong, though we of course do not guarantee that they are the last word. Full hyperparameter details are given in S.I. and in the online codebase.

We note that while these ML methods can attain over 99% accuracy on the full MNIST training set (6000 samples per class, with spatial structure) (LeCun et al., 1995), the few-samples regime is fundamentally different, and standard ML methods are evidently not well-suited to it, compared to biological systems. This few-samples regime requires that (for MNIST) one ignores 99.9% of the usual training data. In addition, the downsampling and vectorization of input images affect accuracy.

Given  $N < 10$  training samples per class, MothNet (with log-likelihood classifier) substantially outperformed all the ML methods, and slightly outperformed ML methods at  $N = 10$  (Fig. 5). However, MothNet (log-likelihood) appears to have limited capacity, and maxed out at a mean accuracy of  $\approx$ 75%. Thus, the ML methods began to pull ahead of MothNet (log-likelihood) at  $N = 30$ –100, depending on ML method.

MothNet (with softmax classifier) had similar accuracy to the NN, outperforming NNs at  $N = 1$  and being roughly equivalent thereafter. The ML methods never pulled ahead of MothNet (softmax) for  $N \leq 100$ . Mean post-training accuracies of the various methods are plotted in Fig. 5, vs number of training samples per class.

*Classifier details.* The MothNet accuracies shown in Fig. 5 are for MothNet instances randomly generated according to templates from Delahunt et al. (2018b), with various learning rates and numbers of “sniffs” (see 4.4 and 4.5 for details). Nearest-Neighbors and SVM used built-in Matlab functions with z-scaled input features, and with other hyperparameters (number-of-neighbors, box constraint) optimized for each  $N$ . The Neural Net was coded in Tensorflow. It had one hidden layer with 85 units (more layers and/or units did not help), Gaussian noise, and dropout. Learning rate and number of epochs were optimized for each  $N$ . For Nearest Neighbors, SVM, and NN, the features were the 85 vectorized pixels from images pre-processed exactly as for MothNet.

The CNN was a packaged Matlab example, with the number of epochs optimized for each  $N$ . Alone of all the methods, the CNN used standard MNIST ( $28 \times 28$ ) images with spatial content, on which it can attain  $>99\%$  accuracy given  $N \geq 500$ . Using full-sized images obviated the need to alter the CNN network parameters to fit smaller images, and thus guaranteed a capable architecture.

MothNet with log-likelihood classifier uses statistics from the validation set, as described in 3.3. MothNet with softmax classifier does not, in common with the ML methods.

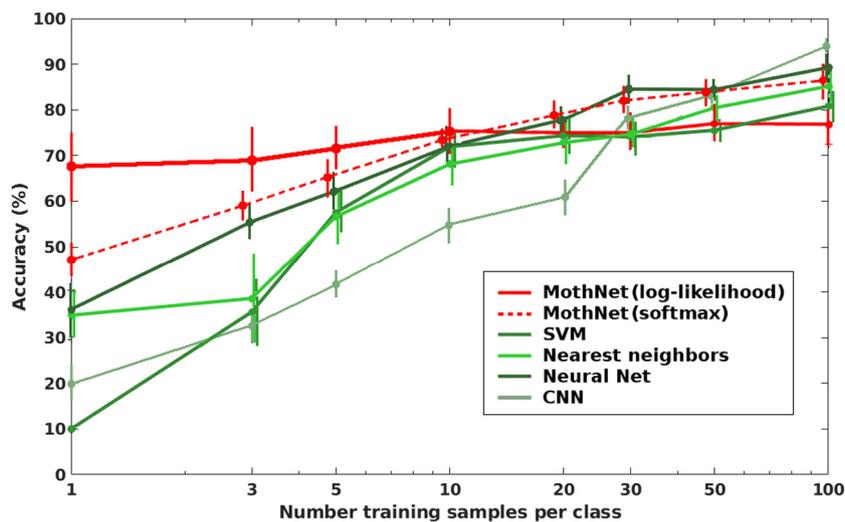
### 4.3. Comparison to ML one-shot methods

A variety of ML methods focus on one-shot learning. These have been tested on Omniglot, an MNIST-like dataset with 1623 alphabetic characters with 20 samples each, created and introduced by Lake, Salakhutdinov, and Tenenbaum (2015). Some of these ML methods were also applied to MNIST.

We applied MothNet to the Omniglot dataset, and where possible compared it with published results from these one-shot ML methods, on both Omniglot and  $\nu$ MNIST.



**Fig. 4.** Statistics of EN response distributions, before and after training, from 13 MothNet instances, showing how training causes the class response distributions to diverge in each EN. The top left subplot shows an EN's naive responses (all ENs had similar responses). The blue arrow represents training. The numbered subplots show the situation post-training: Let  $\mu_{ij}$  be the mean response of the  $i$ th EN to digits of class  $j$  in a single MothNet experiment ( $\mu_{ij}$  will be different pre- and post-training). Let  $\sigma_{ij}$  be the std dev of the  $i$ th EN's responses to digits of class  $j$  in that same MothNet experiment ( $\sigma_{ij}$  will be different pre- and post-training). Then in the  $i$ th numbered subplot, the  $j$ th dot gives the mean( $\mu_{ij}$ ), and the  $j$ th bar gives  $\pm$  mean( $\sigma_{ij}$ ), over post-training  $\mu_{ij}$  and  $\sigma_{ij}$  from 13 MothNet instances. The EN's responses to digits in the targeted class are in blue, and responses to the non-targeted classes are in gray. Mean trained accuracy for this template was 76%, range 71%–83%. This figure is a population-level, statistical analogue to the individual MothNet results shown in Fig. 3. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



**Fig. 5.** Bug vs machine: Mean post-training accuracy vs log of  $N$  = training samples per class. Results are shown for MothNet (in red), as well as for Nearest-Neighbors, SVM, and Neural Nets (all on downsampled  $v$ MNIST, i.e. no spatial info); and CNN (on standard MNIST). At  $N < 10$  training samples per class, MothNet significantly outperformed standard ML methods. Bars show std devs. 13 runs per data point. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

We note that these methods have distinct goals beyond the one-shot learning task. Lake, Salakhutdinov, Gross, and Tenenbaum (2011) seeks insight into human transfer learning via extraction of stroke patterns from drawn characters. Woodward and Finn (2017) explores an active learning system to speed up learning via good training sample choice. MANN (Santoro, Bartunov, Botvinick, Wierstra, & Lillicrap, 2016) augments a NN with an external memory. The Matching Network (Vinyals, Blundell, Lillicrap, Kavukcuoglu, & Wierstra, 2017) also adds external memory to a NN or CNN, and has success on complex datasets (e.g. ImageNet). The Siamese Network (Koch, Zemel, & Salakhutdinov, 2015) uses pairs of NNs or CNNs to develop a similarity detector which, tuned by pre-training, can detect similarities in new class instances. The Neural Statistician (Edwards & Storkey,

2016) learns to compute statistics of new datasets from few samples. (The Ladder Network (Rasmus, Valpola, Honkala, Berglund, & Raiko, 2015) leverages a small labeled training set and a large unlabeled training set from the same classes, a powerful method but for a use-case not treated here.) This diversity of goals makes comparison between methods somewhat beside the point, and does not do justice to the various strengths of the methods. However, the results give evidence that MothNet performance on this one-shot task is comparable to that of specialized one-shot ML methods.

All comparisons were for one-shot learning. Test datasets included vectorized Omniglot with 20 and with 5 test classes, and  $v$ MNIST. Methods that required spatial datasets were omitted. A major difference in the methods is that all the ML methods

**Table 1**  
Comparison of one-shot methods on vectorized datasets.

| Method            | Test data    | # pre-training | Test accuracy (%) |
|-------------------|--------------|----------------|-------------------|
| Siamese NN        | Omniglot 20x | 24k            | 58                |
| MANN              |              | 24k            | 61 (15x)          |
| MothNet (Softmax) |              | 0              | 31                |
| MothNet (logLike) |              | 0              | 65                |
| MANN              | Omniglot 5x  | 24k            | 82                |
| Woodward/Finn     |              | 24k            | 72 (2-shot)       |
| MothNet (softmax) |              | 0              | 57                |
| MothNet (logLike) |              | 0              | 76                |
| Siamese NN        | vMNIST 10x   | 24k            | See text          |
| MothNet (softmax) |              | 0              | 47                |
| MothNet (logLike) |              | 0              | 72                |

Notes: Results for ML methods are drawn from their papers, and are restricted to methods that were tested on vectorized datasets. For ML methods' accuracy on vMNIST, see "Results on vMNIST" in text. Mann Omniglot 20x results are for 15, not 20 classes ("15x"). Woodward/Finn Omniglot 5x results are for two training samples per class ("2-shot"). MothNet results on Omniglot are averages over 11 runs, each with randomly selected target classes.

require pre-training on a large set of samples from separate-but-similar classes (transfer learning), while MothNet assumes no pre-training option.

Due to the larger size of Omniglot images, the number of input features in MothNet's architecture was increased from 85 to 200, with MB size increased proportionally. Because Omniglot images are pure binary, they were smeared with a 2-D Gaussian before vectorization and input to MothNet.

*Results on vectorized Omniglot 20x and 5x.* MothNet (log-likelihood) was similar to ML methods, while MothNet (softmax) trailed all ML methods (Table 1). Thus, while the MothNet model arguably encodes an equally strong classifier, it cannot access its full potential without using a validation set.

*Results on vMNIST.* No non-spatial ML methods reported results on vMNIST. However, there were two clear trends in ML methods' accuracies:

(i) ML methods' performances on spatial MNIST were roughly 20% lower than their performances on spatial Omniglot (Siamese CNN 92→70%; Matching CNN 94→72%; Neural Statistician 93→79%), presumably due to mismatch between the MNIST characters and the Omniglot characters used for pre-training. In contrast, MothNet's accuracy was similar on Omniglot and vMNIST tasks. (ii) ML methods' performances on non-spatial Omniglot were lower than on spatial Omniglot (e.g. Siamese NN vs CNN, 58% vs 92%).

Combining these two trends we estimate that MothNet (softmax) accuracy was similar to, and MothNet (log-likelihood) accuracy exceeded, the projected performance of ML methods on vMNIST. For example, Siamese NN accuracy on vMNIST might be estimated from CNN accuracy on spatial MNIST and the ratio of spatial (CNN) vs vectorized (NN) Siamese accuracies on Omniglot, i.e.  $70(\frac{92}{58}) \approx 44 \approx$  MothNet (softmax), and  $\ll$  MothNet (log likelihood).

To summarize results of one-shot learning on these vectorized datasets:

(i) When the Test set closely matched the pre-training data, MothNet encoded a similar amount of class information as the ML methods, but required a validation set to fully access it.

(ii) When the Test set was somewhat distinct from the pre-training data (i.e. vMNIST vs Omniglot), MothNet encoded more class information and attained similar or higher accuracy.

(iii) MothNet (uniquely) needed no pre-training.

#### 4.4. Growth rate effects, one-shot learning

The speed at which MothNet learns, i.e. the number of training samples required to reach maximum accuracy, is determined to large degree by the Hebbian growth rate, and the related decay

rate, on MB→EN connections. When the connection weights hit the rails of their dynamic range, no further learning is possible (we did not modify the maximum synaptic weight constraints). We ran experiments to study the effect of Hebbian growth rates on post-training accuracy, given various training set sizes. One MothNet template (the "natural learner") had a biologically plausible growth rate, while the second template (the "fast learner") had a growth rate "turned up to 11". High growth rate templates allowed us to test MothNet's skill at one-shot learning (i.e. given just one training sample per class).

Fast learners attained strong immediate accuracy (mean 75%) on just one training sample, but additional training samples provided no further gain. This is because the MB→EN synaptic weights have fixed outer rails: Growth must stop when  $w_{ij}$  hits the top rail, and decay must stop when  $w_{ij} = 0$ .

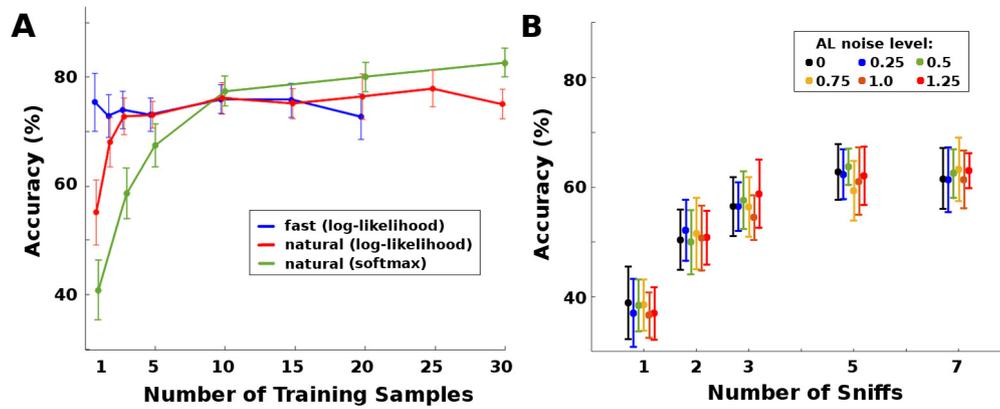
Natural learners took several (~20) training samples to attain maximum accuracy, but that final accuracy was slightly higher. The trade-off of learning speed vs maximum attained accuracy is seen in Fig. 6A. The solid curves show fast- and natural-learner accuracies vs number of training samples.

#### 4.5. Sniffing, effects of AL noise

*Sniffs.* Biological NNs have a remarkable ability to learn from very few training samples. In addition, "sniffing" behavior (repeated sampling of a given odor) is a common biological strategy (Martin et al., 2011; Wilson, 2008). We ran experiments to see whether sniffing behavior improved the various stages of learning. Sniffing applied to test samples did not improve test accuracy (experiments not shown).

However, sniffing during training had a strong effect, especially in one-shot regimes. When a "natural learner" (i.e. with learning rate set to a biologically reasonable level) was given a single training sample (one-shot), multiple sniffs raised post-training accuracy from 35% to over 60%. Five sniffs delivered maximal increase, with no further gains from additional sniffs. This increase in accuracy due to sniffing in the one-shot context is seen in Fig. 6B. The use of sniffing may partly explain why MothNet was able to reach 80% accuracy with fewer samples per class (20 vs 200+) than the honeybee MB model in Huerta and Nowotny (2009).

*AL noise.* The moth AL is a very noisy system. That is, neural responses to a given odor stimulus (or absent any odor) have high variance. To see whether this AL noise was beneficial to learning performance, we also varied the AL noise level in MothNet during the sniffing experiments, from near-zero through the high end of biologically reasonable (per *in vivo* data). We had expected



**Fig. 6. A:** Effects of learning rate parameters. Mean accuracy  $\pm$  std dev ( $\mu \pm \sigma$ ) for a natural learner (red), a fast learner (blue), and a learner using softmax classifier (green). For all data points, number of sniffs was varied to maximize trained accuracy. The fast learner attained 75% accuracy in one-shot, but with no further gains. The natural learner started lower but ultimately attained higher accuracy. There was no difference between fast and natural learners using Softmax. **B:** Effects on one-shot learning of (i) AL noise and (ii) sniffing. (i) AL noise: Each cluster of  $\mu \pm \sigma$  bars represent varying levels of AL noise, with low-to-high noise plotted in different colors from left-to-right at each  $x$ -axis location. AL noise level did not affect accuracy. (ii) Multiple sniffs greatly improved one-shot accuracy. Results for a “natural” learner with log-likelihood accuracies are shown; Softmax results were similar. 13 MothNet instances per data point. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the presence of AL noise to improve performance, given multiple sniffs, by acting as a kind of sample augmentation, analogous to that used in deep NN training. In fact, noise levels in the AL had no effect on accuracy. This is seen in Fig. 6B, where each cluster of mean  $\pm$  std dev bars shows different levels of AL noise for a given number of sniffs.

#### 4.6. Sparsity experiments

High-dimensional, sparse neural layers are a widespread motif in biological NNs (Ganguli & Sompolinsky, 2012). To examine the effects of sparsity in the context of learning, we trained a reasonably capable MothNet template, varying only the sparsity level in the MB during training (17 MothNet instances per sparsity level). Sparsity here is measured as the fraction of MB neurons that are responsive to stimuli (1% is very sparse, 50% is very dense). In MothNet, MB sparsity levels (in both non-training and training modes) are parameters. Sparsity levels in the MB affected two crucial behaviors: Intra-class signal-to-noise ratio (SNR) of EN responses; and how well a given EN $_j$ ’s response to training focused on class  $j$  (learning focus).

Results show that the sparse MB layer plays a key role in learning by controlling and focusing Hebbian weight updates.

High MB response fraction (i.e. high density, or low sparsity) correlated strongly with high SNR (i.e. reliability of intra-class EN responses). But it also resulted in poor post-training classifier accuracy, because Hebbian growth was not focused: Rather, it boosted all weights due to the excess of active MB neurons, so weight increases into a particular EN were not restricted to just the most class-relevant signals. The trained EN then responded strongly to all digits rather than just its target digit.

Conversely, low MB response (i.e. high sparsity) resulted in low intra-class SNR, since not enough MB neurons would fire in response to a digit for reliable activation of the ENs. But high sparsity correlated strongly with high “learning focus”: Training focused gains on the correct class, ensuring that ENs’ responses to their targeted digits were preferentially strengthened. This resulted in stronger post-training accuracy.

Post-training discrimination accuracy appeared to represent a tradeoff between intra-class SNR and learning focus. This tradeoff is shown in Fig. 7, which plots these effects of different sparsity levels in the MB, as they relate to learning:

*Learning focus.* (red curve) This plots a figure-of-merit defined as the average standardized distance between EN response distributions to trained and control classes,

$$LF = \frac{1}{9} \sum_{i \neq j} \frac{(\mu E_{ij} - \mu E_{ji})}{0.5(\sigma E_{ij} + \sigma E_{ji})}, \text{ notation as in Eq. (5).} \quad (6)$$

This is average Bhattacharyya distance if EN response distributions are Gaussian. In very sparse regimes training was strongly focused on the trained class, while in dense regimes training “raised all boats” and the trained response distributions were poorly separated. Thus high sparsity focused learning well, and low sparsity diluted it. This finding matches results in Huerta and Nowotny (2009) and Peng and Chittka (2017).

*Intra-class signal-to-noise ratio.* (black dotted curve) This plots the mean intra-class signal-to-noise ratio (SNR):

$$SNR = \frac{\mu(f)}{\sigma(f)} \text{ where } f = \text{EN odor response;} \quad (7)$$

This was an opposite situation: A very sparse MB resulted in high intra-class variance in trained EN responses (low SNR), while a denser MB delivered much more consistent within-class responses (high SNR).

*Trained accuracy.* (blue curve) This is a fit to post-training classifier accuracies (a natural objective function) from experiment.

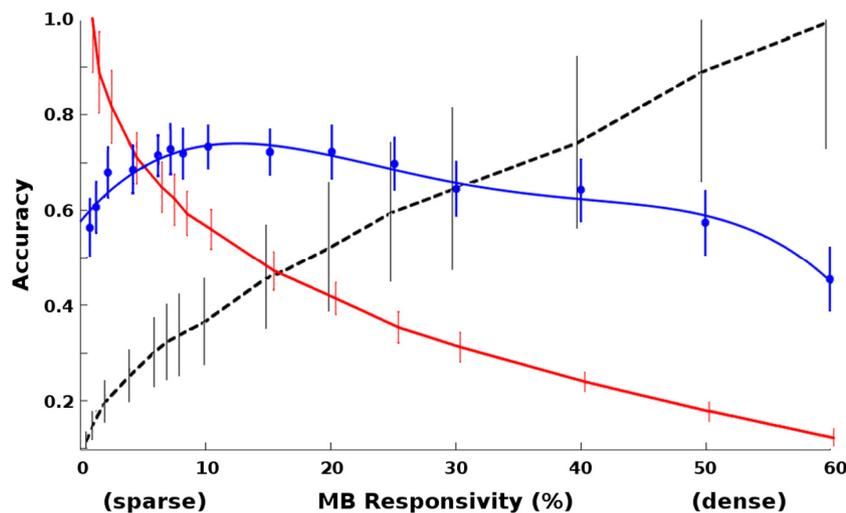
Judged by best post-training classification accuracy, optimal MB sparsity level for MothNet appears to represent a compromise between delivering sufficient intra-class SNR and sufficient “learning focus” for inter-class distinctions. This optimal region lies somewhere between 5% and 20%, consistent with the 5%–15% commonly observed in live BNNs.

## 5. Discussion

### 5.1. A biological toolkit for building NNs

Our goal has been to provide a proof-of-concept as to the learning abilities of even the simplest of biological neural networks (BNNs), and to clarify how a NN built from biological elements learns a predictive model for a general ML task.

In order to learn new odors, the moth olfactory network uses just a few core tools: A noisy pre-processing network with



**Fig. 7.** Effects of MB sparsity on learning: Blue domed curve = a fit to mean trained accuracies, which peaked at 5%–20% sparsity, a compromise between learning focus and high intra-class signal-to-noise ratio (SNR). Means  $\pm$  std devs ( $\mu \pm \sigma$ ) are shown. Red curve =  $\mu \pm \sigma$  of learning focus (separation of responses to trained vs control digits). Black dotted curve =  $\mu \pm \sigma$  intra-class SNR. Learning focus and SNR are scaled to [0, 1] for plotting. 17 MothNet instances per sparsity level.

competitive inhibition; Hebbian plasticity controlled by a high-dimensional sparse layer; and generalized (global) stimulation during training.

Our key finding is that a neural net built with this biological toolkit can succeed at a general rapid-learning task, and in fact can out-perform standard ML methods. MothNet learned to read *v*MNIST digits, increasing its accuracy more than  $5\times$  (from 15% to 75%–80%) given only a few training samples (1–20 per class). Also, because Hebbian weight updates focus entirely on activity induced by the class being trained (not on activity induced by control classes, as in backprop), new classes can be added and trained without retraining on existing classes. Further, the biological tools analyzed here are well-suited to being combined and stacked into larger, deeper neural nets, just as convolutional kernels, maxpool, etc., are combined to build current DNNs.

In these experiments the MothNet model stayed close to the moth’s very simple olfactory architecture, and used a simple feature set as input (pixels). That a NN based on a moth, i.e. a bottom rung on the ladder of biological neural complexity, can match or exceed current ML methods on some problems argues for the potential value of applying BNN-based tools in the ML context. The success of live BNNs at a wide range of tasks further argues for this potential. We believe that a free hand with network design, such as using better input features, biologically-unrealistic model parameters, and more complex architectures, would yield strong results in a variety of tasks.

### 5.2. The rapid learning regime

It is not, on reflection, surprising that an insect brain might out-perform ML methods at a few-samples learning task. First, BNNs in general excel at rapid learning from few samples. Second, in the typical ML context, such as a competition or a comparison with other methods using well-defined benchmarks, the number of training samples is fixed and is often high (e.g. 60,000 for MNIST). The competitive pressure is to maximize accuracy, with no penalty for using lots of training data. In sharp contrast, an insect pays a high cost for every extra training datum, and their competitive pressure is to very rapidly attain “good-enough” accuracy.

These two regimes complement each other, and there are applications where it may be advantageous to pair a fast, rough learner with a slower, more precise learner. One example might

be an adaptive controls learner for a drone: If the drone collides in mid-air and loses an engine, you eventually want to learn a new, optimal control strategy for the reduced system. But the first necessity is to learn a “good enough” control strategy very fast, before crashing, to allow time for the subtler system to train. Another example might be a multi-stage bootstrap labeling scheme when labeled data are scarce: In the initial stage, given just one or a few labeled samples, the fast learner might classify and apply labels to more samples (perhaps keeping only those with the most certain labels) until there are enough labeled samples to train a subtler system as the second stage labeler.

### 5.3. Role of the Hebbian update rule

We hypothesize that the ability of BNNs such as MothNet to generalize well from very few samples is related to the Hebbian update mechanism, which operates differently than statistically-based optimizers. It has no objective function or output-based loss that is pushed back through the network as in backprop or agent-based reinforcement learning. Rather, weight updates occur on a purely local “use it or lose it” basis. Typical ML optimizers, such as backprop, excel at interpolation (Mallat, 2016), but arguably not extrapolation. Perhaps Hebbian update rules, which strengthen or decay connection weights according to amount of traffic during learning episodes, enable effective generalization outside the convex hull of interpolation.

### 5.4. Role of the sparse layer

A second finding emphasizes the key role of sparse layers during learning. Sparse, high-dimensional layers are a widespread motif in biological neural systems, especially in networks related to memory and plasticity (Ganguli & Sompolinsky, 2012). In the MothNet model, the sparse layer (MB) plays a vital role in learning because all the plastic synapses connect into or out of the sparse layer, allowing it to modulate the Hebbian updates to the synaptic connections by taking advantage of the fact that Hebbian growth is an AND gate (“fire-together, wire-together”) (Delahunt et al., 2018b). This ensures that learning boosts the important signal (i.e. the signal associated with a given sample) and not artifacts.

While the sparse MB layer calls to mind the backprop sparse autoencoder (Ng, 2010), the biological role described here has no

obvious analogue in backprop sparse encoders, since it is tied to the Hebbian update method. However, biological sparse layers may also perform functions similar to those found in backprop sparse autoencoders, such as reducing noise (Vincent, Larochelle, Bengio, & Manzagol, 2008) and reducing the dimension of the feature space in order to better match the essential dimension of the classification task (Makhzani & Frey, 2013).

### 5.5. Role of competitive inhibition

A key feature of the AL layer is competitive inhibition between units, each of which receives input from exactly one feature. This competitive inhibition may enable rapid learning by creating several attractor basins for inputs, each focused on a particular class according to which features present most strongly. This might serve to push otherwise similar samples (of different classes) away from each other, towards their respective class attractors, increasing the effective distance between the samples. Thus the outputs of the AL, after this competitive inhibition, may have better separation by class. The fact that the MothNet model includes an AL may be one reason it was able to learn more rapidly than the honeybee MB-only model in Huerta and Nowotny (2009).

### 5.6. Role of octopamine

Unanswered by these experiments is whether generalized stimulation by octopamine is required in ML systems such as MothNet, distinct from actual biological systems. In the moth, octopamine stimulation may offer a work-around to avoid biological constraints on Hebbian growth rates and input intensity. That is, octopamine may act primarily as an accelerant. However, engineered NNs can easily crank up growth rates, enforce higher MB activity, and amplify signals during training, all without recourse to the octopamine mechanism but perhaps with the same beneficial effect on learning. In this case, generalized stimulation would not be a necessary part of the biological toolkit described here for application to ML tasks. However, it may be that octopamine stimulation also enables exploration of the coding solution space not normally activated by stimuli. Alternate ways to replace this functionality in the ML context are not so obvious.

### 5.7. Role of noise

Also unclear is the role during learning, if any, of the high intrinsic noise built into MothNet's Antennal Lobe. In a biological context neural noise (including in the AL) can have diverse benefits independent of learning (Delahunt, Maia, & Kutz, 2018; Knight, 1972; Ma, Beck, Latham, & Pouget, 2006; Wiesenfeld & Moss, 1995). However, our experiments found no positive reason to build high intrinsic noise into the AL (or any NN layer) in an ML context. However, a noisy pre-amp (e.g. the AL) might still be directly beneficial to learning, for three reasons.

First, when coupled with sniffing, a noisy AL might provide a version of data augmentation (as used in DNNs) by distorting the codes delivered to the MB and readout neurons. This would improve one-shot or few-shot learning, if the distortions induced by the noisy AL somewhat mimicked the within-class sample variation (this condition was likely absent in our MNIST experiments). Second, injecting noise into input layers, or corrupting training samples, can improve NN classification performance (An, 1996), suggesting a concrete benefit during training (though perhaps not during classification). Third, injecting noise may be a useful or even necessary way to explore the solution space (Bengio & Fischer, 2015).

### 5.8. Expanding the biological toolkit

The algorithmic tools for learning characterized here derive from one of the simplest biological learning networks (a bug brain). It is still unclear whether modifications to this network can substantially improve performance, or whether this particular architecture can achieve just so much in terms of accuracy and scope of work, commensurate with the constrained needs and pressures experienced by insects as they evolved. Happily, BNNs are incredibly diverse and complex. Countless other useful elements can likely be abstracted from other biological systems and applied to ML tasks. This requires accurately modeling these more complex systems and analyzing how they learn, an open-ended task.

### Acknowledgments

Thanks to:

Christian Szegedy for valuable draft-editing input;  
Satpreet Singh for insights about the role of the AL;  
Adam P. Jones for his work on the Python version of MothNet;  
The referees, for highly germane feedback and insights.

CBD's work was partially supported by the Swartz Foundation, United States.

JNK acknowledges support from the Air Force Office of Scientific Research, United States (FA9550-19-1-0011).

### Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.neunet.2019.05.012>.

### References

- An, G. (1996). The effects of adding noise during backpropagation training on a generalization performance. *Neural Computation*, 8(3), 643–674. URL <http://dx.doi.org/10.1162/neco.1996.8.3.643>.
- Babadi, B., & Sompolinsky, H. (2014). Sparseness and expansion in sensory representations. *Neuron*, 83(5), 1213–1226. URL <http://www.sciencedirect.com/science/article/pii/S0896627314006461>.
- Bazhenov, M., & Stopfer, M. (2010). Forward and back: Motifs of inhibition in olfactory processing. *Neuron*, 67(3), 357–358. URL <http://www.sciencedirect.com/science/article/pii/S0896627310005842>.
- Bengio, Y., & Fischer, A. (2015). Early inference in energy-based models approximates back-propagation. arXiv e-prints, abs/1510.02777. URL <https://arxiv.org/abs/1510.02777>.
- Bhandawat, V., Olsen, S. R., Gouwens, N. W., Schlieff, M. L., & Wilson, R. I. (2007). Sensory processing in the *Drosophila* antennal lobe increases reliability and separability of ensemble odor representations. *Nature Neuroscience*, 10, 1474–1482.
- Campbell, R., Honegger, K., Qin, H., Li, W., Demir, E., & Turner, G. (2013). Imaging a population code for odor identity in the *Drosophila* mushroom body. *Journal of Neuroscience*, 33(25), 10568–10581.
- Campbell, R. A., & Turner, G. C. (2010). The mushroom body. *Current Biology*, 20(1), R11 – R12. URL <http://www.sciencedirect.com/science/article/pii/S096098220901851X>.
- Caron, S., Ruta, V., Abbott, L., & Axel, R. (2013). Random convergence of olfactory inputs in the *Drosophila* mushroom body. *Nature*, 497(5), 113–117.
- Cassenaer, S., & Laurent, G. (2007). Hebbian STDP in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. *Nature*, 448, 709 EP –. URL <http://dx.doi.org/10.1038/nature05973>.
- Dacks, A. M., Christensen, T. A., Agricola, H. J., Wollweber, L., & Hildebrand, J. G. (2005). Octopamine-immunoreactive neurons in the brain and subesophageal ganglion of the hawkmoth *Manduca sexta*. *Journal of Comparative Neurology*, 488(3), 255–268. 15952164[pmid]. URL <https://www.ncbi.nlm.nih.gov/pubmed/15952164>.
- Dacks, A., Riffell, J., Martin, J., Gage, S., & Nighorn, A. (2012). Olfactory modulation by dopamine in the context of aversive learning. *Journal of Neurophysiology*, 108(2), 539–550.
- Dayan, P., & Abbott, L. F. (2005). *Theoretical neuroscience: computational and mathematical modeling of neural systems*. The MIT Press.
- Delahunt, C. B., Maia, P. D., & Kutz, J. N. (2018). Built to last: Functional and structural mechanisms in the moth olfactory network mitigate effects of neural injury. arXiv. URL <https://arxiv.org/abs/1808.01279>.

- Delahunt, C. B., Riffell, J. A., & Kutz, J. N. (2018b). Biological mechanisms for learning: A computational model of olfactory learning in the *Manduca sexta* moth, with applications to neural nets. *Frontiers in Computational Neuroscience*, 12, 102, URL <https://doi.org/10.3389/fncom.2018.00102>.
- Edwards, H., & Storkey, A. (2016). Towards a neural statistician. arXiv. URL <https://arxiv.org/abs/1606.02185>.
- Fukushima, K. (1980). Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36, 193–202.
- Galizia, C. G. (2014). Olfactory coding in the insect brain: data and conjectures. *European Journal of Neuroscience*, 39(11), 1784–1795.
- Ganguli, S., & Sompolinsky, H. (2012). Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Annual Review of Neuroscience*, 35(1), 485–508, PMID: 22483042.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. Vol. 1. MIT press Cambridge.
- Hammer, M., & Menzel, R. (1995). Learning and memory in the honeybee. *Journal of Neuroscience*, 15(3), 1617–1630, URL <http://www.jneurosci.org/content/15/3/1617>.
- Hammer, M., & Menzel, R. (1998). Multiple sites of associative odor learning as revealed by local brain microinjections of octopamine in honeybees. *Learning and Memory*, 5(1), 146–156, 10454379[pmid]. URL <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC311245/>.
- Hebb, D. O. (1949). *The organization of behavior: a neuropsychological theory* (p. xix, 335). Wiley New York.
- Hige, T., Aso, Y., Rubin, G. M., & Turner, G. C. (2015). Plasticity-driven individualization of olfactory coding in mushroom body output neurons. *Nature*, 526, 258 EP –. URL <http://dx.doi.org/10.1038/nature15396>.
- Higham, D. J. (2001). An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Review*, 43(3), 525–546, URL <http://dx.doi.org/10.1137/S0036144500378302>.
- Honegger, K. S., Campbell, R. A. A., & Turner, G. C. (2011). Cellular-resolution population imaging reveals robust sparse coding in the *Drosophila* mushroom body. *Journal of Neuroscience*, 31(33), 11772–11785, URL <http://www.jneurosci.org/content/31/33/11772>.
- Hong, E. J., & Wilson, R. I. (2015). Simultaneous encoding of odors by channels with diverse sensitivity to inhibition. *Neuron*, 85(3), 573–589, URL <http://www.sciencedirect.com/science/article/pii/S0896627314011490>.
- Hubel, D., & Wiesel, T. (1962). Receptive fields, binocular interaction, and functional architecture in the cat's visual cortex. *Journal of Physiology*, 160, 106–154.
- Huerta, R., & Nowotny, T. (2009). Fast and robust learning by reinforcement signals: Explorations in the insect brain. *Neural Computation*, 21(8), 2123–2151.
- Knight, B. W. (1972). Dynamics of encoding in a population of neurons. *The Journal of General Physiology*, 59(6), 734–766, URL <http://jgp.rupress.org/content/59/6/734>.
- Koch, G., Zemel, R., & Salakhutdinov, R. (2015). *Siamese neural networks for one-shot image recognition*. URL <http://www.cs.utoronto.ca/~gkoch/files/msc-thesis.pdf>.
- Kvello, P., Löfaldli, B., Rybak, J., Menzel, R., & Mustaparta, H. (2009). Digital, three-dimensional average shaped atlas of the *Heliothis virescens* brain with integrated gustatory and olfactory neurons. *Frontiers in Systems Neuroscience*, 3, 14, URL <https://www.frontiersin.org/article/10.3389/neuro.06.014.2009>.
- Lake, B., Salakhutdinov, R., Gross, J., & Tenenbaum, J. (2011). One shot learning of simple visual concepts. In *Proceedings of the 33rd annual conference of the cognitive science society*.
- Lake, B., Salakhutdinov, R., & Tenenbaum, J. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266), 1332–1338.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436.
- LeCun, Y., & Cortes, C. (2010). *MNIST handwritten digit database*. Website. URL <http://yann.lecun.com/exdb/mnist/>.
- LeCun, Y., Jackel, L., Bottou, L., Brunot, A., Cortes, C., Denker, J., et al. (1995). Comparison of learning algorithms for handwritten digit recognition. In *International conference on artificial neural networks* (pp. 53–60).
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9, 1432 EP –. Article. URL <http://dx.doi.org/10.1038/nn1790>.
- Makhzani, A., & Frey, B. J. (2013). K-sparse autoencoders. *CoRR*, abs/1312.5663, URL <http://arxiv.org/abs/1312.5663>.
- Mallat, S. (2016). Understanding deep convolutional networks. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 374(2065), URL <http://rsta.royalsocietypublishing.org/content/374/2065/20150203>.
- Martin, J. P., Beyerlein, A., Dacks, A. M., Reisenman, C. E., Riffell, J. A., Lei, H., et al. (2011). The neurobiology of insect olfaction: Sensory processing in a comparative context. *Progress in Neurobiology*, 95(3), 427–447, URL <http://www.sciencedirect.com/science/article/pii/S0304008211001742>.
- Masse, N. Y., Turner, G. C., & Jefferis, G. S. (2009). Olfactory information processing in *Drosophila*. *Current Biology*, 19(16), R700 – R713, URL <http://www.sciencedirect.com/science/article/pii/S0960982209013013>.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4), 115–133, URL <https://doi.org/10.1007/BF02478259>.
- Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. The MIT Press.
- Ng, A. (2010). *Sparse autoencoder*. Webpage. URL <https://web.stanford.edu/class/archive/cs/cs294a/cs294a.1104/sparseAutoencoder.pdf>.
- Peng, F., & Chittka, L. (2017). A simple computational model of the bee mushroom body can explain seemingly complex forms of olfactory learning and memory. *Current Biology*, 27(2), 224–230, URL <http://www.sciencedirect.com/science/article/pii/S096098221631288X>.
- Perisse, E., Burke, C., Huetteroth, W., & Waddell, S. (2013). Shocking revelations and saccharin sweetness in the study of *Drosophila* olfactory memory. *Current Biology (London)*, 23(17), R752–R763, S0960-9822(13)00921-4[PII], 24028959[pmid]. URL <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3770896/>.
- Rasmus, A., Valpola, H., Honkala, M., Berglund, M., & Raiko, T. (2015). Semi-supervised learning with ladder networks. arXiv. URL <https://arxiv.org/abs/1507.02672>.
- Riffell, J. A., Lei, H., Abrell, L., & Hildebrand, J. G. (2012). Neural basis of a pollinator's buffet: Olfactory specialization and learning in *Manduca sexta*. *Science*, URL <http://science.sciencemag.org/content/early/2012/12/05/science.1225483>.
- Roelfsema, P. R., & Holtmaat, A. (2018). Control of synaptic plasticity in deep cortical networks. *Nature Reviews Neuroscience*, 19, 166 EP –. Review article. URL <http://dx.doi.org/10.1038/nrn.2018.6>.
- Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65–386.
- Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., & Lillicrap, T. P. (2016). One-shot learning with memory-augmented neural networks. *CoRR*, abs/1605.06065, URL <http://arxiv.org/abs/1605.06065>.
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61(Supplement C), 85–117, URL <http://www.sciencedirect.com/science/article/pii/S0893608014002135>.
- Shields, V., & Hildebrand, J. (2001). Responses of a population of antennal olfactory receptor cells in the female moth *Manduca sexta* to plant-associated volatile organic compounds. *Journal of Comparative Physiology A*, 186(12), 1135–1151, URL <https://doi.org/10.1007/s003590000165>.
- Vincent, P., Larochelle, H., Bengio, Y., & Manzagol, P. A. (2008). Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on machine learning ICML '08*, (pp. 1096–1103). New York, NY, USA: ACM, URL <http://doi.acm.org/10.1145/1390156.1390294>.
- Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., & Wierstra, D. (2017). Matching networks for one shot learning. arXiv. URL <https://arxiv.org/abs/1606.04080>.
- Wiesenfeld, K., & Moss, F. (1995). Stochastic resonance and the benefits of noise: from ice ages to crayfish and SQUIDS. *Nature*, 373, 33.
- Wilson, R. I. (2008). Neural and behavioral mechanisms of olfactory perception. *Current Opinion in Neurobiology*, 18(4), 408–412, Sensory systems. URL <http://www.sciencedirect.com/science/article/pii/S0959438808000883>.
- Woodward, M., & Finn, C. (2017). Active one-shot learning. arXiv. URL <https://arxiv.org/abs/1702.06559>.