Neural Networks 125 (2020) 56–69

Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

# Neuromodulated attention and goal-driven perception in uncertain domains

Xinyun Zou [a,*], Soheil Kolouri [b], Praveen K. Pilly [b], Jeffrey L. Krichmar [c,a]

[a] Department of Computer Science, University of California, Irvine, Irvine, CA 92697, USA
[b] Information and Systems Sciences Laboratory, HRL Laboratories LLC, Malibu, CA 90265, USA
[c] Department of Cognitive Sciences, University of California, Irvine, Irvine, CA 92697, USA

ARTICLE INFO

*Article history:*
Received 7 August 2019
Received in revised form 13 December 2019
Accepted 27 January 2020
Available online 1 February 2020

*Keywords:*
Neuromodulation
Goal-driven perception
Uncertainty
Top-down attention
Contrastive excitation backprop

ABSTRACT

In uncertain domains, the goals are often unknown and need to be predicted by the organism or system. In this paper, contrastive Excitation Backprop (c-EB) was used in two goal-driven perception tasks – one with pairs of noisy MNIST digits and the other with a robot in an action-based attention scenario. The first task included attending to even, odd, low, and high digits, whereas the second task included action goals, such as "eat", "work-on-computer", "read", and "say-hi" that led to attention to objects associated with those actions. The system needed to increase attention to target items and decrease attention to distractor items and background noise. Because the valid goal was unknown, an online learning model based on the cholinergic and noradrenergic neuromodulatory systems was used to predict a noisy goal (*expected uncertainty*) and re-adapt when the goal changed (*unexpected uncertainty*). This neurobiologically plausible model demonstrates how neuromodulatory systems can predict goals in uncertain domains and how attentional mechanisms can enhance the perception for that goal.



## 1. Introduction

Artificial attentional mechanisms in neural networks tend to respond to sensory inputs similarly regardless of context and goals (Itti & Koch, 2000; Tsotsos, Eckstein, & Landy, 2015; Zhang et al., 2018). However, biological systems select relevant information to guide behavior in the face of noisy and unreliable signals, as well as rapidly adapt to unforeseen situations. Goal-driven perception treats the same situation differently based on context and effectively directs attention to goal-relevant inputs. Often, these goals are unknown and must be learned through experience. Moreover, these goals or contexts can shift without warning. Goal-driven perception helps prevent overemphasis on less relevant stimuli and instead focus on critical stimuli that require an immediate response.

In the brain, neuromodulators are important contributors to attention and goal-driven perception. In particular, the cholinergic (ACh) system drives bottom-up, stimulus-driven attention, as well as top-down, goal-driven attention (Avery, Dutt, & Krichmar, 2014). Furthermore, the ACh system increases attention to task-relevant stimuli, while decreasing attention to distractions (Baxter & Chiba, 1999; Oros, Chiba, Nitz, & Krichmar, 2014). This procedure is similar to the core idea behind contrastive Excitation Backprop (c-EB). In c-EB, a top-down excitation mask increments attention to the target features, and an inhibitory mask decrements attention to distractors (Zhang et al., 2018). The noradrenergic (NE) system responds to surprises or large deviations from priors (Yu & Dayan, 2005). When the NE system responds phasically, where the neural activity rapidly and transiently increases, it causes a network to reset (e.g., by re-initializing activities) that allows rapid adaptation under unseen/new conditions (Bouret & Sara, 2005; Grella et al., 2019).

We modified the c-EB network for use in a goal-driven perception task, where the system needed to increase attention to the intended goal object and decrease attention to the distractor. In the first experiment, we presented pairs of noisy MNIST digits to the neural network. One goal class was to attend to the digit based on its parity (i.e., even or odd goal), and another goal class was to attend based on the magnitude of the digit (i.e., low- or high-value goal). In addition, we added a neuromodulatory model to the head of the network architecture that regulated goal selection. Similar to the model of the ACh and NE neuromodulatory systems proposed by Yu and Dayan (2005), we framed the task as an attentional task where the goal (even, odd, low or high value) needed to be learned from experience (*goal identity*) and the goal might be noisy and rewarded with some probability (*goal validity*). In the second experiment, we generalized our model to an action-based attention scenario, where "eat", "work-on-computer", "read", and "say-hi" were goal actions and the robot

* Correspondence to: 2232 Social & Behavioral Sciences Gateway, University of California, Irvine, Irvine, CA 92697, USA.
*E-mail addresses:* xinyunz5@uci.edu (X. Zou), skolouri@hrl.com (S. Kolouri), pkpilly@hrl.com (P.K. Pilly), jkrichma@uci.edu (J.L. Krichmar).

https://doi.org/10.1016/j.neunet.2020.01.031
0893-6080/

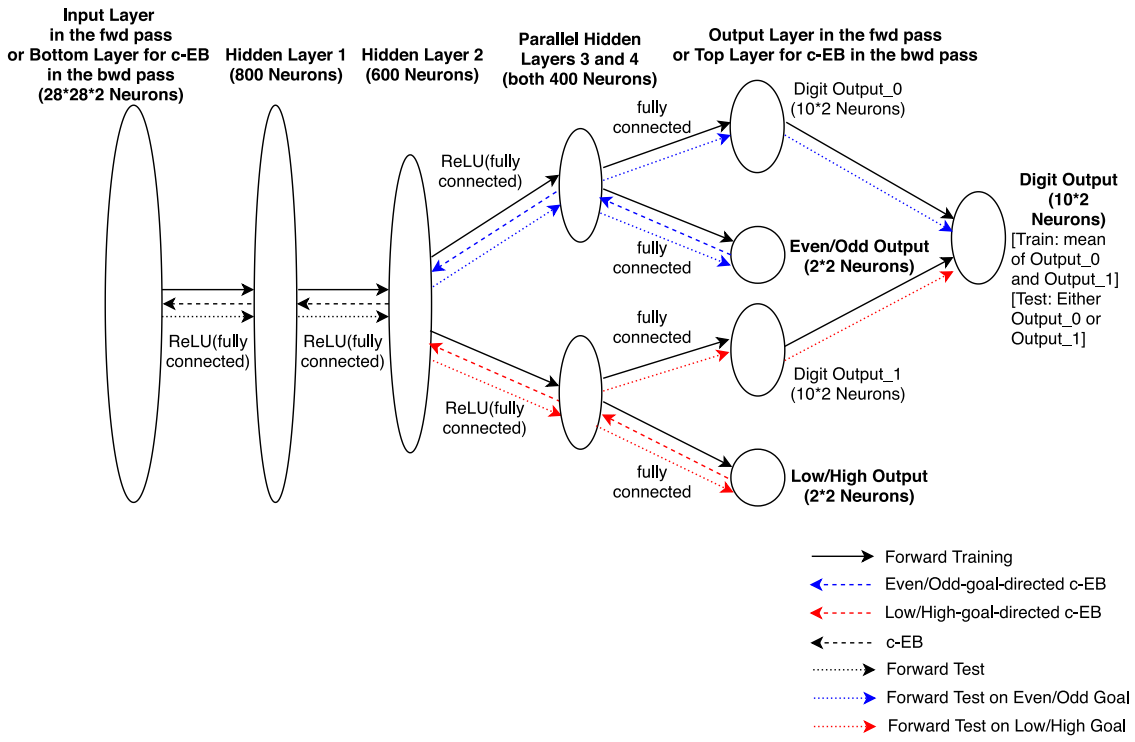


**Fig. 1.** Network setup for our bottom-up classification process and our top-down attentional search process, with a pair of noisy MNIST digits as the input data in the forward pass. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

needed to attend to and retrieve objects that corresponded to the action.

## 2. Methods

We modified the c-EB neural network model to attend to different goals. Section 2.1 describes how we tested the ability of the network to increase attention to different goals and digits for the noisy MNIST-pair experiment. Section 2.2 introduces our neuromodulatory learning system to predict unknown and uncertain goals based on experience, still using the noisy MNIST-pair experiment as an example. Section 2.3 describes how our method was generalized to demonstrate goal-driven perception in a human support robot.

### 2.1. Network architecture

Fig. 1 shows our bottom-up classification process and our top-down attentional search process. In the forward pass, the input layer received a pair of 28 × 28-pixel noisy MNIST digits and thus had $28 \times 28 \times 2 = 1568$ neurons (LeCun, Bottou, Bengio, & Haffner, 1998). To test the network's ability to filter out distractions, noise that was randomly set between 0 and 0.7 was added to normalized pixel values (between 0 and 1) of the original MNIST digits. The final pixel values were then normalized again between 0 and 1.

Following the input layer were two sequential fully connected hidden layers with 800 and 600 neurons, respectively. Next, there were two parallel fully connected hidden layers, each with 400 neurons. All neurons in these layers implemented a Rectified Linear Unit (ReLU) as the activation function (Nair & Hinton, 2010). Each of the two parallel hidden layers led to the output in one goal class (parity/magnitude) along with the digit output. For each (left/right) side of the input image, after the third hidden layer, there were two parity (an even and an odd) output neurons and ten digit output neurons, which contributed to the parity/digit prediction using winner-take-all (WTA) on the activation probability of each parity/digit output neuron. Similarly the magnitude/digit prediction was obtained after the fourth hidden layer for each side of the input image. During training, the final digit output took the average of the digit output generated from the two parallel hidden layers. During testing, the final digit prediction was the digit output generated from one of the two parallel hidden layers, depending on the cued goal class.

In our top-down attentional search process (see backward arrows in Fig. 1), one of the four goals (even, odd, low, and high) was selected, which excited the corresponding goal neuron for each of the two digits in the image and inhibited all the other goal neurons at the top layer of the backward pass (i.e., the output layer of the forward pass). The weights were backpropagated from the top layer to one of the parallel hidden layers below to excite the neurons corresponding to the goal (see dashed arrows from the top layer to parallel hidden layers 3 and 4 in Fig. 1). Then the weights at the top layer were converted from excitatory to inhibitory in order to create a mask (note that the weights were originally non-negative). This inhibitory mask was used in an additional backpropagation from the top layer to the parallel hidden layer below corresponding to the goal. The result of a subtraction between the two backpropagations was a contrastive signal (Zhang et al., 2018). This contrastive signal was then used to perform regular EB over the remaining layers, which finally generated the probability of each given pixel in the input layer for exciting the cued goal neurons. In addition to exciting the goal neurons and inhibiting non-goal neurons, the contrastive signal canceled out common winner neurons. Such a contrastive extension of the backpropagation could effectively ignore noisy distractors and lead to more accurate attention focus on the goal (Zhang et al., 2018).

#### 2.1.1. Modification of c-EB

Excitation Backprop (EB) was developed as a goal-driven attentional framework for a CNN classifier based on a probabilistic

winner-take-all (WTA) process (Zhang et al., 2018). It could visualize the features at each layer in the hierarchy that were relevant to a given output neuron. An important extension of EB was to have contrastive Excitation Backprop (c-EB), which discriminated the goal pixels from distractors by canceling out common winner neurons for different goals and amplifying discriminative neurons for the target goal (Zhang et al., 2018).

The EB mechanism kept non-negative weights between activation neurons and used these excitatory connections to transmit top-down signals. The top-down relevance of a neuron $a_n$ in the layer $L_l$ was defined by its probability of being chosen as a layer-wise winner, which was called the Marginal Winning Probability (MWP) $P(a_n)$ (Zhang et al., 2018):

$$P(a_n) = \sum_{a_m \in (L_0, L_1, \ldots, L_{l-1})} (P(a_n|a_m) \cdot P(a_m)), \tag{1}$$

where $a_m$ denoted each parent neuron in the preceding layer(s). The winner neurons were recursively sampled in the top-down direction according to the conditional winning probability $P(a_n|a_m)$ (Zhang et al., 2018):

$$P(a_n|a_m) = \begin{cases} \dfrac{\widehat{a_n} \cdot w_{nm}}{\sum_{n:w_{nm}\geq 0} (\widehat{a_n} \cdot w_{nm})} & \text{if } w_{nm} \geq 0, \\ 0 & \text{otherwise,} \end{cases} \tag{2}$$

where $w_{nm}$ was the weight between a parent neuron $a_m$ and one child neuron $a_n$, and $\widehat{a_n}$ denoted a non-negative activation response.

For c-EB, the contrastive signals were transmitted in the top-down fashion to obtain highly discriminative attention maps (i.e., contrastive MWP maps) in the target layer (i.e., the bottom layer in Fig. 1). Extending from Eq. (1), Zhang et al. (2018) defined the contrastive MWP (c-MWP) of the target layer $L_l$ as

$$A - \overline{A} = P_0 \cdot (P_1 - \overline{P_1}) \cdot P_1 \cdot \ldots \cdot P_{l-1}, \tag{3}$$

where $A$ represented the MWP, $\overline{A}$ was the dual MWP for the contrastive units, $P_0$ was the signal from the guessed goal, and $\overline{P_1}$ was the conditional probability of the inhibition mask from the top layer. The weights for the inhibition mask were the negation of the original weights from the top layer. Therefore, the threshold condition for $\overline{P_1}$ was the reverse of that for $P_1$ in Eq. (2).

We extended the PyTorch (Paszke et al., 2017) implementation of c-EB (Greydanus, 2018), whereas the original code for c-EB (Zhang et al., 2018) was written in Caffe (Jia et al., 2014). Different from their implementation, our network included different goal classes labeled for each of the two noisy MNIST digits. In addition, the c-EB was processed through one of the two parallel hidden layers immediately below the top layer in the backward pass of our network.

The system increased attention to the digit corresponding to the selected goal and decreased attention to the distractor digit. One goal class attended to the digit based on its parity (i.e., either odd or even), whereas the other goal class attended to the digit based on its magnitude (i.e., low values between 0 and 4 inclusively or high values between 5 and 9 inclusively). This resulted in two goals within each goal class. After supervised training on noisy pairs generated from the MNIST training dataset, c-EB was applied to the top-down attentional process on the test pairs and driven by one of the four goals to excite only the pixels relevant to the goal digit.

Fig. 2 shows the two noisy test pairs and their c-EB generated attention maps according to each goal. c-EB driven by a goal went through the backward pass and excited the pixel neurons only related to the goal digit. On the irrelevant digit side, most pixel neurons were inhibited instead. Furthermore, background noise on both sides were ignored. Therefore, the goal digits and goal

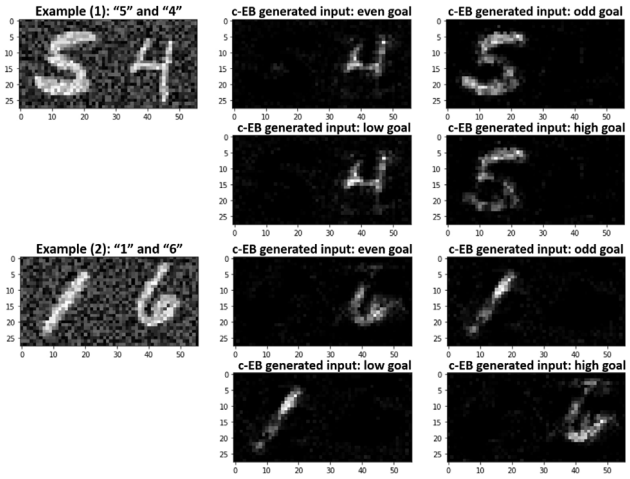


**Fig. 2.** Two example test pairs of noisy MNIST digits and their c-EB highlighted results. The two digits in each test pair had the opposite goals both in the parity (even/odd) goal class and in the magnitude (low/high) goal class. These restrictions were not applied to the training pairs during the experiments.

identity neurons could all be predicted correctly with high certainty in the end of the forward pass in these examples. However, even if two goal identities targeted the same goal digit, their highlighted pixels in c-EB visualization were not all the same. It is reasonable because our model had different output heads for the different goal classes.

With the first noisy test pair of "5" and "4" in Fig. 2 as an example, if an even goal was selected, then both even goal neurons in the left and right digit sides were activated and all other goal neurons were inhibited in the top layer of the backward pass. After the c-EB process that sent contrastive signals through the third hidden layer and added them up for normal EB via the second and first hidden layers to reach the bottom layer (see Fig. 1), pixels related to the digit "4" were highlighted, whereas pixels related to the distractor digit "5" and the background noise were inhibited.

### 2.1.2. Training and testing process

The training process consisted of incremental learning on noisy MNIST pairs. The original MNIST dataset was split into 60,000 digit images for training and 10,000 digit images for testing (LeCun et al., 1998). At each training step, 256 pairs of noisy MNIST digits were randomly selected and modified from the original MNIST training set. Every 200 training steps, 2000 noisy MNIST pairs, randomly selected and modified from the original MNIST test set, were used for validation to evaluate the current training progress. The training process stopped after 4400 steps, when the validation accuracy was the highest. Therefore, a total of 1,126,400 noisy MNIST pairs were used for training. The following test process consisted of 10,000 pairs randomly selected and modified from the original MNIST test set. Given the sizes of the original MNIST training and test sets, there must be digit overlap within some training pairs or within some test pairs. However, there was no digit overlap between training and test. The two digits in each training pair could have the same or opposite parity (even/odd) and the same or opposite magnitude (low/high), whereas those in each test pair all had the opposite parity and the opposite high and low values.

During training, a log-softmax function, followed by a negative log likelihood function, was applied after the forward pass to the neurons in the output layers that represented a digit, even parity, odd parity, low value, or high value. Then the sum of loss was used to calculate the gradient for each parameter in

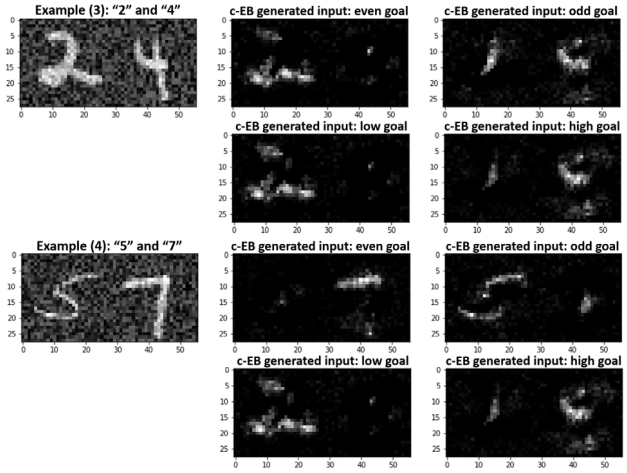


**Fig. 3.** Two more example test pairs of noisy MNIST digits and their c-EB highlighted results. The two digits in each test pair had the same goals both in the parity (even/odd) goal class and in the magnitude (low/high) goal class. The tested condition with same parity and/or high/low was not included in later experiments.

the model. At the end of each training step, a parameter update was performed based on the current gradient calculated using the Adam optimizer (Kingma & Ba, 2014) with a learning rate of 0.001.

During testing, c-EB drove goal-driven perception by increasing the activity of input neurons corresponding to the goal digit and masking out the neurons corresponding to the distractor digit. In the top layer of our network (Figs. 1 and 4), either two out of the four parity neurons or two out of the four magnitude neurons were activated, depending on the selected goal. For example, if an "odd" goal was selected, the odd neuron for the left digit and the odd neuron for the right digit were both excited, whereas all other goal neurons for both digits were inhibited. This resulted in c-EB in the backward pass to increase attention to this goal and its corresponding digit, and to ignore all other goals (see blue arrows for a parity goal and red arrows for a magnitude goal in Fig. 1). Such c-EB generated input, with pixels highlighted for the goal digit (Fig. 2), then went through the forward pass again in a way similar to the training process. However, according to the goal identity, only the parallel hidden layer related to the target goal class was used to predict the goal digit.

If a test pair had same parity or same high or low values (see Fig. 3), an existing goal would drive c-EB to highlight pixels from both digits. As expected, it shows that this ambiguous situation would cause confusion in the attention system. We assume it would cause confusion and random selection by a human facing the same stimuli. Thus, we did not present same parity or same high or low values as test pairs.

After this training and testing procedure, the parameters of the fully trained model were fixed for the neuromodulated goal-driven perception experiments.

### 2.2. Neuromodulated goal-driven perception

The overall neuromodulated procedure of goal-driven perception is shown in Fig. 4. As described in Section 2.1.2, the network was trained with pairs of noisy MNIST digits to learn the digits and their parity (even/odd-value) and magnitude (low/high-value) goal classes. Then it was tested by selecting one of the even, odd, low-value, and high-value goals to trigger c-EB in the backward pass and generate an attention map, which further led to prediction of the digit and goal in the succeeding forward pass. After the robustness of c-EB prediction was verified, we applied ACh and NE neuromodulatory neurons to track the expected and unexpected uncertainties respectively and guess the goal for each trial. The guessed goal was applied to the top layer for c-EB as the intended goal for the current test pair. The guessed goal and the predicted digit were compared with the true goal and the true goal-related digit. The prediction was used to modify the neuromodulatory activities for the next trial, as will be described in Section 2.2.1.

#### 2.2.1. ACh and NE neuromodulation

**Algorithm 1** ACh and NE Neuromodulation Process

**Constant Input:** $\beta = 0.7$, $num\_switches = 10$,
$K = 4$, $trial\_interval = 400$, $trial\_range = 30$,
$ne^{reset} = 0.25$, $ne^{min} = 0.25$, $ne^{max} = 1.0$,
$ch^{reset} = 1.0$, $ch^{min} = 0$, $ch^{max} = 10.0$,
$ne_{correct} = 0.70$, $ne_{wrong} = 1.10$,
$ch_{correct} = 1.40$, $ch_{wrong} = 0.90$
**Other Input:** *all_test_pairs*, *validity_options*
Initialize $ACh_i$ to $ch^{reset}$ for $i = 1, 2, ..., K$.
Initialize $NE$ to $ne^{reset}$.
Set *minLen* to ($trial\_interval - trial\_range$).
Set *maxLen* to ($trial\_interval + trial\_range$).
**for** $q = 1$ **to** $num\_switches$ **do**
  Randomly set *majorGoal* from 0, 1, ..., K-1.
  Set *minorGoal* from the same goal class.
  Randomly set *validity* from *validity_options*.
  Randomly set *trialLen* within [*minLen*, *maxLen*].
  **for** $t = 1$ **to** *trialLen* **do**
    Pick a new *test_pair* from *all_pairs*.
    Randomly set $r$ between [0, 1.0).
    **if** $r <$ *validity* **then**
      Set *trueGoal* to *majorGoal*.
    **else**
      Set *trueGoal* to *minorGoal*.
    **end if**
    Select *guessGoal* from Softmax (see Eq. (4)).
    Get *trueDigit* from *test_pair* with *trueGoal*.
    Apply *guessGoal* to the top layer.
    Obtain *map* via c-EB (see Eq. (3)).
    Get *predDigit* via fwd pass with *map*.
    Compare *predDigit* with *trueDigit*.
    Compare *trueGoal* with *guessGoal*.
    Update ACh and NE (see Eqs. (5) and (6)).
    Compute the reset threshold $\theta^{reset}$ (see Eq. (7)).
    **if** $NE > \theta^{reset}$ **then**
      Reset $ACh_i$ to $ch^{reset}$ for $i = 1, 2, ..., K$.
      Reset $NE$ to $ne^{reset}$.
    **end if**
  **end for**
**end for**

For goal-driven perception, the network must select a goal when they are uncertain and unknown *a priori*. Similar to a model of the ACh and NE neuromodulatory systems proposed by Yu and Dayan (2005), the goal target (even, odd, low, or high value) was rewarded with a probability (*goal validity*), but that goal would change periodically (*goal identity*). ACh neurons tracked *expected uncertainties* of the potential goals. NE neurons tracked *unexpected uncertainties*, and responded phasically when a goal identity change was detected. When the NE system responded phasically, it caused a network reset by re-initializing the ACh and NE neural activities, which allowed rapid adaptation under novel conditions.

Algorithm 1 shows the logic of our ACh and NE neuromodulatory model. There were $K = 4$ ACh neurons, each corresponding

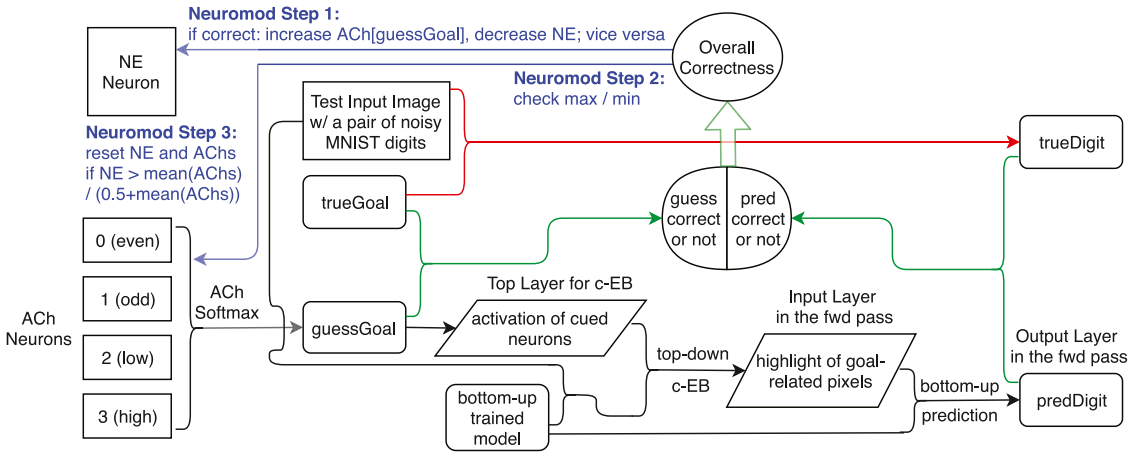


**Fig. 4.** The neuromodulated procedure of making digit prediction from a guessed goal for the noisy MNIST-pair experiment.

to a goal (i.e., even, odd, low, or high value), and one NE neuron. One of the four attentional goal tasks was selected as the major goal for each goal switch. The true goal identity in each trial was set to either the major goal or the minor goal according to the goal validity (see Section 2.2.2 for details). The true goal digit was obtained from the labels of the test pair by using the true goal identity. The activities of ACh neurons were input to a softmax function for goal selection:

$$p(goal)_i = \frac{\exp(\beta \cdot ACh_i)}{\sum_{j=1}^{K} \exp(\beta \cdot ACh_j)} \text{ for } i = 1, 2, \ldots, K, \tag{4}$$

where $\beta$ was the temperature governing exploration versus exploitation and $p(goal)_i$ was the probability of selecting goal $i$. This guessed goal activated two neurons related to the goal in the top layer of our network architecture (Figs. 1, 4), which directed c-EB in the backward pass to activate the goal-relevant pixels in the test pair and then predicted the digit in the forward pass. If the prediction was correct (which means that the guessed goal identity matched the true goal identity and the predicted digit matched the true goal digit), the ACh level corresponding to the guessed goal (i.e., $ACh_g$) increased and the NE level decreased; the opposite would happen otherwise:

$$(ACh_g)_t = \begin{cases} \min(ch_{correct}(ACh_g)_{t-1}, ch^{max}) & \text{if correct,} \\ \max(ch_{wrong}(ACh_g)_{t-1}, ch^{min}) & \text{otherwise,} \end{cases} \tag{5}$$

$$NE_t = \begin{cases} \max(ne_{correct} \cdot NE_{t-1}, ne^{min}) & \text{if correct,} \\ \min(ne_{wrong} \cdot NE_{t-1}, ne^{max}) & \text{otherwise,} \end{cases} \tag{6}$$

where $[ch^{min}, ch^{max}]$ and $[ne^{min}, ne^{max}]$ were ranges for ACh and NE levels. $ch_{correct}$ and $ne_{wrong}$ must be set within $[1.0, 2.0)$, and $ch_{wrong}$ and $ne_{correct}$ must be within $(0, 1.0]$. If the NE level was above a threshold $\theta^{reset}$, ACh and NE activities were reset to baseline levels (Yu & Dayan, 2005):

$$\theta^{reset} = \frac{\left(\sum_{i=1}^{K} ACh_i\right)/K}{0.5 + \left(\sum_{i=1}^{K} ACh_i\right)/K}. \tag{7}$$

Our settings for constant parameters of the neuromodulation process are listed in Algorithm 1. However, there was a wide range of parameter values that could be used to produce stable results. The randomness in Algorithm 1 followed a uniform distribution within the ranges specified, except that selecting *guessGoal* required the softmax distribution.

### 2.2.2. Goal selection

We added an online neuromodulatory model (Fig. 4 and Algorithm 1) to the head of the network architecture in the backward pass to regulate goal selection automatically. In these experiments, the goal (with goal identities of even, odd, low, or high value) needed to be learned from experience. It might be noisy and rewarded with some probability (i.e., goal validity).

Automatic goal selection was tested in 10 runs to measure the average performance. In each run, one of the four attentional goal tasks was randomly selected as the major goal, which stayed the same every 400 ± 30 trials for 10 switches. The minor goal identity came from the same goal class as the major goal identity. For example, if the major goal was "high", then the minor goal became "low" in the same magnitude goal class; or if the major goal was "even", then the minor goal became "odd" in the same parity goal class. The true goal identity was set to either the major goal or the minor goal randomly according to the validity distribution per trial. The true goal digit was obtained from the labels of the test pair of noisy MNIST digits using the true goal identity.

The goal validity values (i.e., 0.99, 0.85, and 0.70) were chosen to correspond with Yu and Dayan (2005). The major goal validity was randomly chosen among the three values each time the major goal identity got switched in a run. The minor goal validity was (1 − major goal validity).

## 2.3. Action-based attention in a robot experiment

To test whether the goal-driven perception model could generalize to a more real-world application, we tested the model in an action-based attention task on the Toyota Human Support Robot (HSR) (Yamamoto, Nishino, Kajima, Ohta, & Ikeda, 2018). For the second experiment, we had four goal actions (i.e., "eat", "work-on-computer", "read", "say-hi") that were associated with images of objects seen by the HSR. Given a desired action, the task for the HSR was to guess the action and direct attention to the object in the scene that could achieve that action. For example, the action "eat" might result in attention to an "apple".

For object classification, we used the Microsoft COCO dataset (Lin et al., 2014) to train a GoogLeNet (Szegedy et al., 2015) via the Caffe framework (Jia et al., 2014) instead of the MNIST-pair network shown in Fig. 1. An advantage of the COCO dataset was it used segmentation to localize individual object instances in an image, which was more accurate and more helpful for top-down attention than using bounding boxes.

For each run of this experiment, desired actions were randomly switched every 50 trials for 10 switches. In each trial,

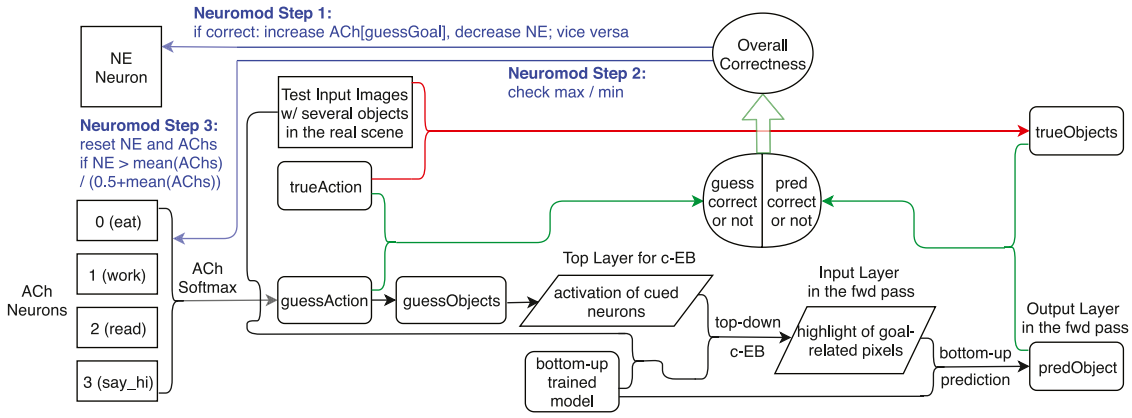


**Fig. 5.** The neuromodulated procedure of making object prediction from a guessed goal action for the indoor robot experiment.

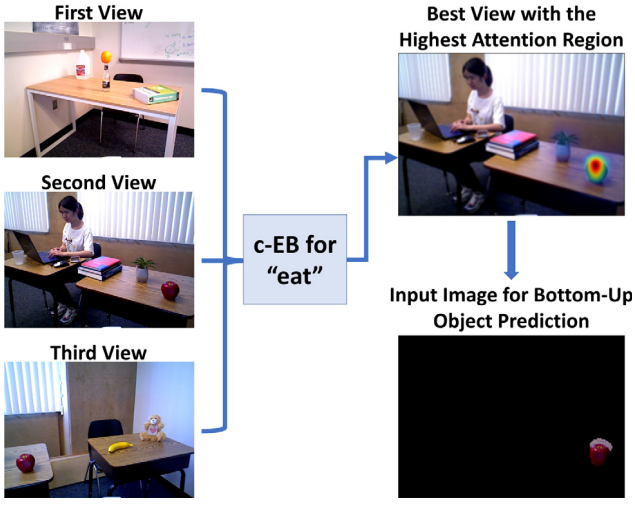


**Fig. 6.** The top-down attentional search process for a guessed action "eat" based on three different real indoor views to select the highest attention region for bottom-up object prediction.

**Fig. 7.** The test scenario for the indoor robot experiment.

each image from three capture angles within an indoor scene was loaded as the input to the pretrained bottom-up model and went through a forward pass with the output layer specified as "loss3/classifier" in Caffe. The number of output prediction classes was set to 80, same as the number of object labels available for the COCO dataset (Lin et al., 2014). Then the c-EB method was applied for the top-down attention process. As in the MNIST-pair experiment, the ACh–NE neuromodulation process with softmax

**Table 1**
Relationship between goal actions and object labels.

| Action | Role | Objects in role |
|---|---|---|
| eat | obj | banana, apple, sandwich, orange, donut, carrot, broccoli, hot dog, pizza, cake |
| | instr | fork, knife, spoon, bowl, cup |
| work | instr | laptop, tv, mouse, keyboard |
| read | obj | book |
| | instr | laptop, cell phone |
| say-hi | obj | person |

on the ACh activities was applied for goal (action) selection (see Fig. 5).

Table 1 linked the guessed action with all related objects – from 80 COCO labels regardless of their semantic roles – which might or might not exist in the test scenario. Those activated object neurons in the top layer drove the c-EB through the second top layer "pool5/7x7_s1" and then normal EB to the bottom layer "pool3/3x3_s2" to generate attention maps related to this guessed goal action. The notation "pool5/7x7_s1" referred to a pooling layer with a kernel size of $7 \times 7$ and a stride of 1. Only the highest attention region (with normalized attentional strength above the threshold of 0.1) corresponding to one of the three real captures maintained its original pixel values, whereas all other parts of the image became black (see Fig. 6). This attention-modulated image became the input to the forward pass of the network and generated object prediction via the top layer. Three conditions needed to be satisfied to generate a overall correct match for that trial: (1) the guessed action matched the true action; (2) the predicted object matched the real object in the scene; (3) the predicted object was associated with the guessed action.

The test environment for this experiment was a classroom scenario as shown in Fig. 7. The test agent was a Toyota HSR (Yamamoto et al., 2018). During each trial, the HSR first guessed an action using the activity of the neuromodulatory neurons and linked it with objects using the semantic network. The HSR then moved from a starting point to the center of the testing scenario, where it captured three images from different view angles using the RGB-D camera. After the attention network predicted an object as described above, the HSR moved towards the object and either picked up the object if it was grabbable (e.g., apple) or pointed at the object with its arm (e.g., laptop). At the trial end, the HSR would present the object to a user who would respond with a "YES" if it matched his/her desired action or otherwise with a "NO". This feedback would be used by the neuromodulatory model to adjust the activity levels of both the

**Table 2**
Prediction for 10,000 test pairs of noisy MNIST digits.

| Goal task | % Correct digit prediction | % Correct goal prediction |
|---|---|---|
| Even | 92.03 | 99.50 |
| Odd | 91.15 | 99.75 |
| Low | 95.39 | 99.54 |
| High | 87.46 | 98.22 |

NE neuron and the ACh neuron related to the current guessed action before guessing the action for the following trial.

## 3. Results

Section 3.1 shows how the network attended to correct goals and predicted the digits corresponding to those goals. Section 3.2 demonstrates the ability of the neuromodulatory head to learn goals based on its experience in uncertain domains. Section 3.3 shows the necessity of having both NE and ACh neurons to correctly predict goals. Section 3.4 compares the performance of our method with two benchmarks. The experiments in these sections were carried out with noisy MNIST pairs. Section 3.5 shows how our goal-driven perception method generalized to an action-based attention task with a physical robot.

### 3.1. Digit prediction with c-EB and noisy MNIST pairs

The training process was carried out for 4400 steps, including 256 noisy MNIST pairs modified from the original MNIST training set per step. The prediction performance of the fully trained model was tested on 10,000 pairs of noisy MNIST digits modified from the original MNIST test set (LeCun et al., 1998). Table 2 shows the digit and goal prediction results with c-EB driven by one of the four goal tasks (i.e., even, odd, low, or high value).

The goal was predicted along with the digit in the output layers for each forward pass. As shown in Table 2, the model predicted the goal correctly over 99% of the time, meaning that after the backward and forward passes the most active neuron for goal prediction matched the true goal. The model predicted the goal digit correctly over 90% of the time, meaning that the most active digit neuron matched the expected digit corresponding to the goal (see Table 2). This indicates that the goal tasks were successfully understood by the c-EB process to highlight related pixels. Although the statistics of the high-value goal task was slightly weaker than that of the other three goal tasks, the performance was still robust overall.

In the next section, we show how this network can autonomously predict goals in uncertain domains.

### 3.2. Goal-driven perception with uncertainties

The robust digit and goal prediction results using c-EB (see Section 3.1) assured that the network architecture could be applied to situations where goals are uncertain and contexts are unknown. Therefore, the next step was to test the reliability and flexibility of our proposed neuromodulation model for predicting goals in a noisy, dynamic environment.

Fig. 8 shows typical runs of our neuromodulated system for three major validity settings. For each major validity of 0.99 (Fig. 8a), 0.85 (Fig. 8b), or 0.70 (Fig. 8c), the first subplot includes the true goals (labeled as "major goal" and "minor goal") and ACh-guessed goals (labeled as "guess"); the second and third subplots show NE and ACh levels. Note that the activity level of the ACh neuron corresponding to a major goal quickly increased, driving attention to the most likely goal as well as suppressing attention to distractors. In cases where the major goal validity was low, the ACh neuron corresponding to the minor goal was also activated, resulting in more exploration and a higher chance of guessing the minor goal. Interestingly, the prediction during exploration tended to remain in the same goal class. When there was a change in the goal identity, the NE neuron quickly recognized the change and responded with spiking activity. This caused the activities in the goal prediction network to reset, and a short period of exploration before the system found the new goal identity. Lower major goal validity led to longer exploration, especially after a goal identity switch, as well as more frequent NE bursts.

We also ran experiments where the goal validity could change during the run. Fig. 9 shows the performance of a typical run with random switching among three major goal validity options, 0.99, 0.85, and 0.70. Similar to Fig. 8, the system in this setting still focused more on the major goal. With lower major goal validity (i.e., when the major goal appeared less frequently, see 0.70 in Fig. 9), the NE neuron fired phasically more frequently; meanwhile, the activity level of the major goal's ACh neuron oscillated more frequently with larger amplitude, giving higher potential for the minor goal's ACh neuron to fire at a low level. Because both the goal validity and goal identity changed during a run, the exploration period lasted longer with a lower major goal validity of 0.85 or 0.70. However, this also led to higher prediction accuracy of the minor goal.

Table 3 shows the performance of goal and digit prediction on the noisy MNIST pairs over 10 runs for each validity setting. The third and fourth columns refer to the percent of trials at which the goal digit prediction was correct. The fifth column refers to the percent of incorrect goal guesses based on the ACh softmax distribution (see Eq. (4), with $\beta = 0.7$). The sixth column refers to the percent of incorrect digit predictions with c-EB driven by the guessed goal, when the ACh-guessed goal already matched the true goal. The seventh column refers to lag length of choosing the correct goals, which was computed as the number of trials between the first trial of a major goal switch and when the network started consistently making correct goal prediction at least 80% of the time over the last 10 trials. The first three rows provide average statistics for runs at which a single goal validity was tested (see also Fig. 8), and the last row corresponds with runs at which the goal validity could change randomly among three options (see also Fig. 9).

### 3.3. Ablation studies

We wanted to understand the effect of each neuromodulator on the network's capability of goal selection. Therefore, we simulated ablation studies on ACh and/or NE neurons in a randomly changing goal validity experiment. These ablations had drastic effects on performance (see Table 4 and Fig. 10) compared to the complete network. With ablation of the NE neuron (see Fig. 10a), there was no scheme for network reset. The ACh neurons were still able to track the major goal switches. However as time elapsed, it took longer for the ACh activity level corresponding to the major goal to rise significantly and properly after each major goal switch, as measured by the lag length. With ablation of the ACh neurons (see Fig. 10b), the goal guessing became random. The firing rate of the NE neuron increased rapidly in the beginning and stayed at extremely high values afterwards. With ablation of both ACh and NE neurons (see Fig. 10c), there was no firing activity of either the NE or ACh neurons, and thus the goal guessing was random. These ablation studies demonstrate the necessity of having one system track the expected uncertainties of goals (ACh) and another respond appropriately when the goal distribution changes (NE).

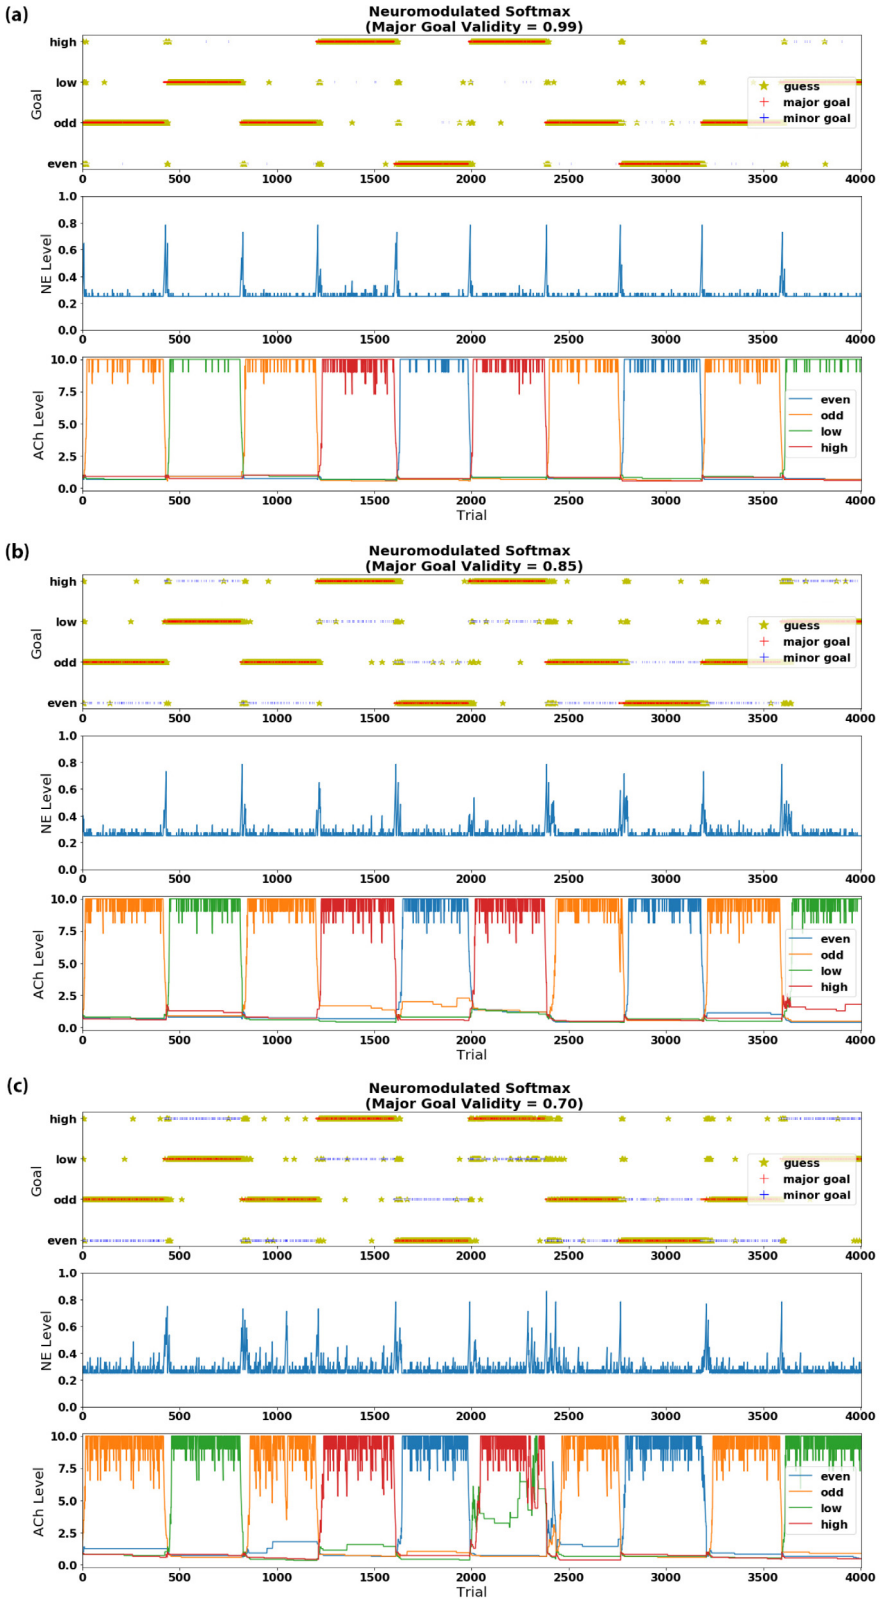


**Fig. 8.** Visualization of goal-driven perception performance on noisy MNIST pairs with the major goal validity chosen from (a) 0.99, (b) 0.85, or (c) 0.70. The major goal identity was randomly picked every 400 $\pm$ 30 trials for 10 switches in a run. The minor goal was the other goal in the same class of the major goal. For example, if the major goal "odd" had validity of 0.70, the minor goal "even" had validity of 0.30 until the next major goal switch. For each major goal validity, the top subplot shows guessed goal identities (in yellow) and true goal identities (either major goals in red or minor goals in blue); the middle and bottom subplots show NE and ACh levels. A softmax function (see Eq. (4), with $\beta = 0.7$) was applied to ACh levels for goal guessing. See text for details. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

**Table 3**
Average goal-driven perception performance on noisy MNIST pairs over 10 runs for each of the four goal validity settings. The first three rows of data correspond with the first experiment of one major goal validity. The last row relates to the second experiment of randomly switched goal validity. *p_valid* means the major goal validity, and $(1 - p\_valid)$ means the minor goal validity. In each run, the major goal was randomly picked every 400 ± 30 trials for 10 switches. The minor goal was selected from the same goal class. The $\beta$ value for the softmax function (see Eq. (4)) was set to 0.7.

| Major goal validity | Minor goal validity | % Correct major goal | % Correct minor goal | % Incorrect ACh softmax goal guessing | % Incorrect c-EB digit prediction | Lag length (trials) |
|---|---|---|---|---|---|---|
| 0.99 | 0.01 | 86.1 | 0.0 | 7.8 | 6.1 | 21 |
| 0.85 | 0.15 | 73.0 | 0.3 | 20.4 | 6.3 | 29 |
| 0.70 | 0.30 | 57.9 | 1.5 | 34.3 | 6.3 | 48 |
| *p_valid* | $1 - p\_valid$ | 75.1 | 0.7 | 18.0 | 6.2 | 30 |

**Table 4**
Average goal-driven perception performance on noisy MNIST pairs over 10 runs for each of the four ablation conditions on the NE and/or ACh neuron(s). In each run, the major goal was randomly picked among the four goal options every 400 ± 30 trials for 10 switches. For each major goal switch, the major goal validity was selected randomly among 0.99, 0.85, and 0.70. The minor goal was selected from the same goal class. The $\beta$ value for the softmax function (see Eq. (4)) was set to 0.7.

| Ablated neuron(s) | % Correct major goal | % Correct minor goal | % Incorrect ACh softmax goal guessing | % Incorrect c-EB digit prediction | Lag length (trials) |
|---|---|---|---|---|---|
| None | 75.1 | 0.7 | 18.0 | 6.2 | 30 |
| NE | 70.8 | 1.1 | 21.6 | 6.5 | 54 |
| ACh | 19.8 | 2.9 | 70.9 | 6.4 | 400 |
| NE & ACh | 19.9 | 2.9 | 70.9 | 6.3 | 400 |

**Table 5**
Average goal-driven perception performance on noisy MNIST pairs over 10 runs among neuromodulated softmax (with $\beta = 0.7$), neuromodulated WTA, and "random-or-fixed" methods. In each run, the major goal was randomly picked among the four goal options every 400 ± 30 trials for 10 switches. For each major goal switch, the major goal validity was selected randomly among 0.99, 0.85, and 0.70. The minor goal was selected from the same goal class.

| Goal-guessing method | % Correct major goal | % Correct minor goal | % Incorrect goal guessing | % Incorrect c-EB digit prediction | Lag length (trials) |
|---|---|---|---|---|---|
| Neuromodulated softmax | 75.1 | 0.7 | 18.0 | 6.2 | 30 |
| Neuromodulated WTA | 76.4 | 0.8 | 16.5 | 6.3 | 24 |
| "Random-or-fixed" | 63.1 | 1.7 | 29.0 | 6.2 | 23 |

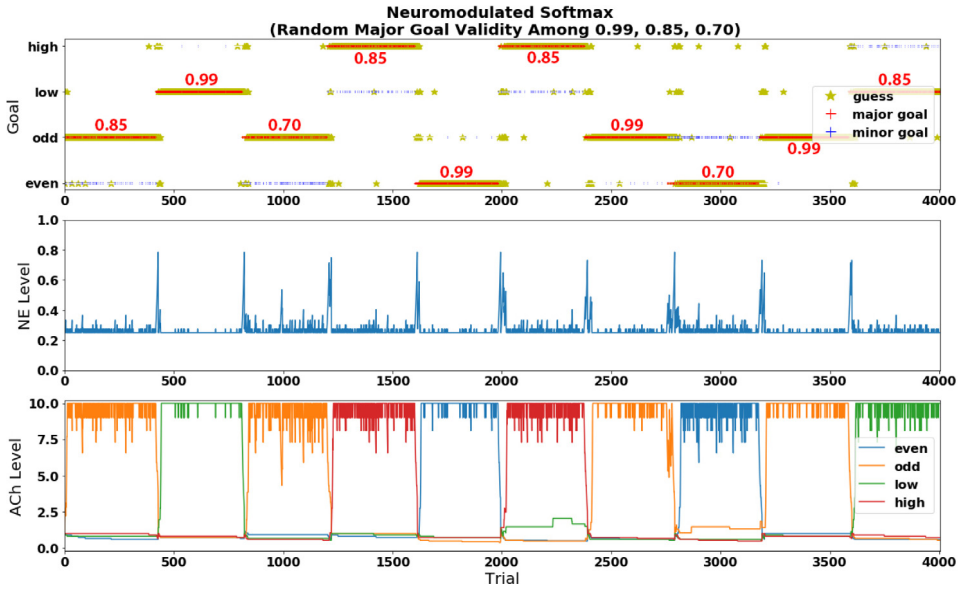


**Fig. 9.** Visualization of goal-driven perception performance on noisy MNIST pairs with the major goal validity randomly switching among 0.99, 0.85, and 0.70. All other settings were the same as shown in Fig. 8. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

### 3.4. Goal selection method comparison

In the neuromodulated procedure of our model (see Fig. 4), the goal was selected by calculating the softmax distribution based on the activities of the four ACh neurons (see Eq. (4), with $\beta = 0.7$). The softmax function was important for raising the chance of choosing the minor goal when the major goal validity was low (e.g., 70%). We compared softmax to a winner-take-all (WTA) selection method, which still used the neuromodulatory head.

In addition, we compared the neuromodulatory head to another benchmark, which we call "random-or-fixed". In the "random-or-fixed" benchmark, a predicted goal was randomly selected until it matched the true goal. Then the goal selection was fixed until a mismatch appeared. In other words, whether the guessed goal was random or stayed the same depended on whether there was a mismatch or match in the goal guessing process of the previous trial.

Table 5 shows the performance comparison among neuromodulated softmax (shown in Fig. 9), neuromodulated WTA (shown

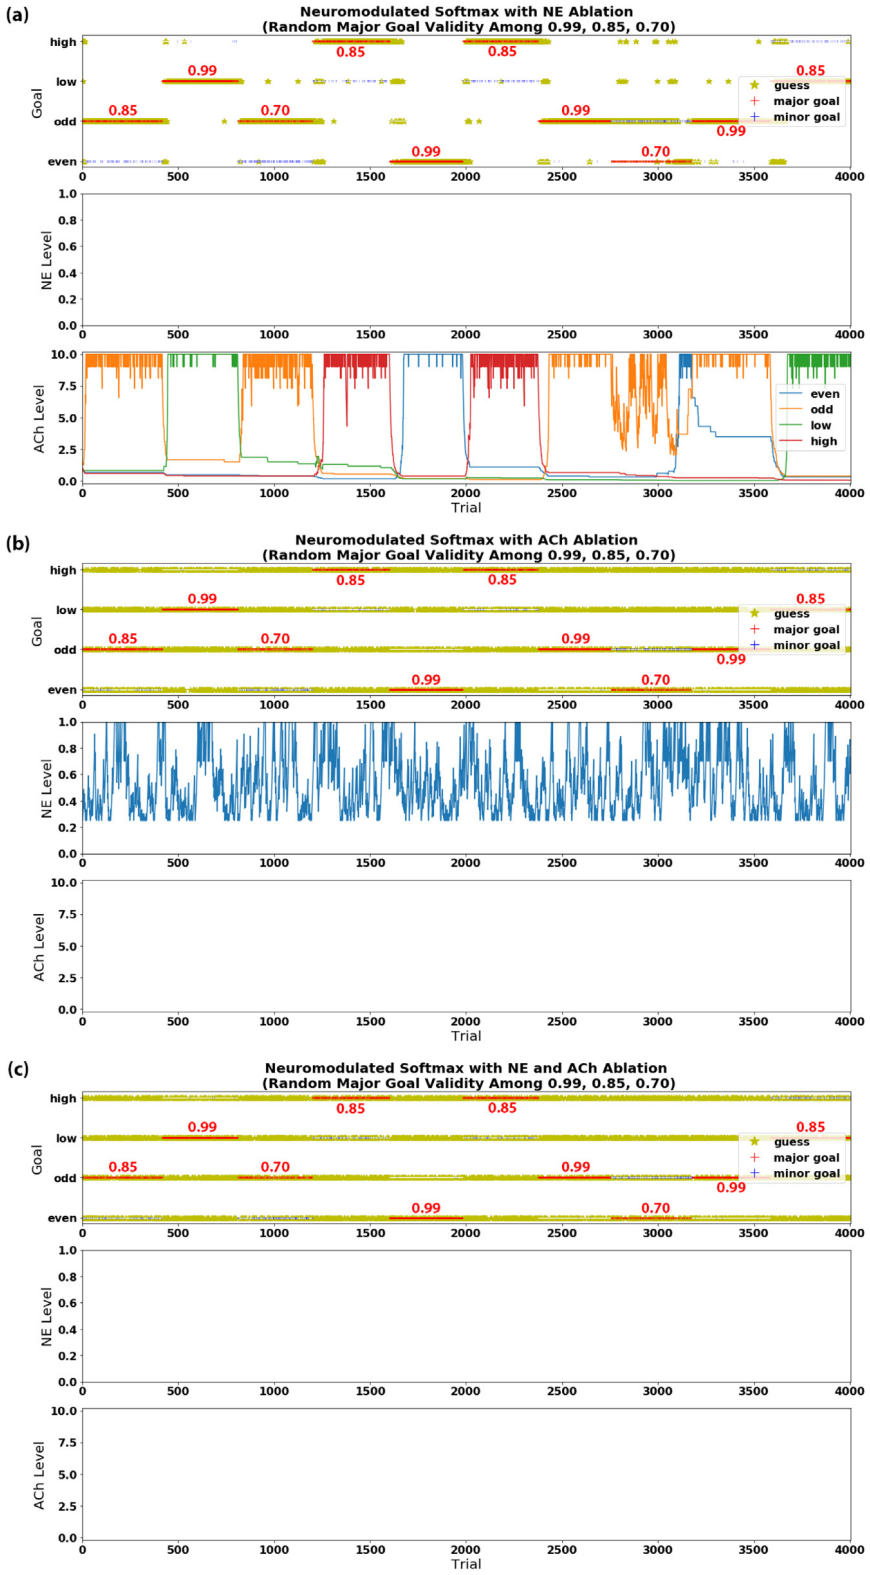


**Fig. 10.** Visualization of goal-driven perception performance on noisy MNIST pairs with the major goal validity randomly switching among 0.99, 0.85, and 0.70, after (a) NE ablation, (b) ACh ablation, and (c) NE and ACh ablation. All other settings were the same as shown in Figs. 8 and 9. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

in Fig. 11), and "random-or-fixed" (shown in Fig. 12) on the noisy MNIST pairs. All three methods had similar lag lengths. Although the "random-or-fixed" method generated the highest percentage of minor goal matches, it was mostly caused by random guesses among all four goals, which also lowered the percentage of major goal matches. Therefore, the neuromodulation process was important for quickly following the desired goal class without hesitating over all four goals after each major goal switch. For

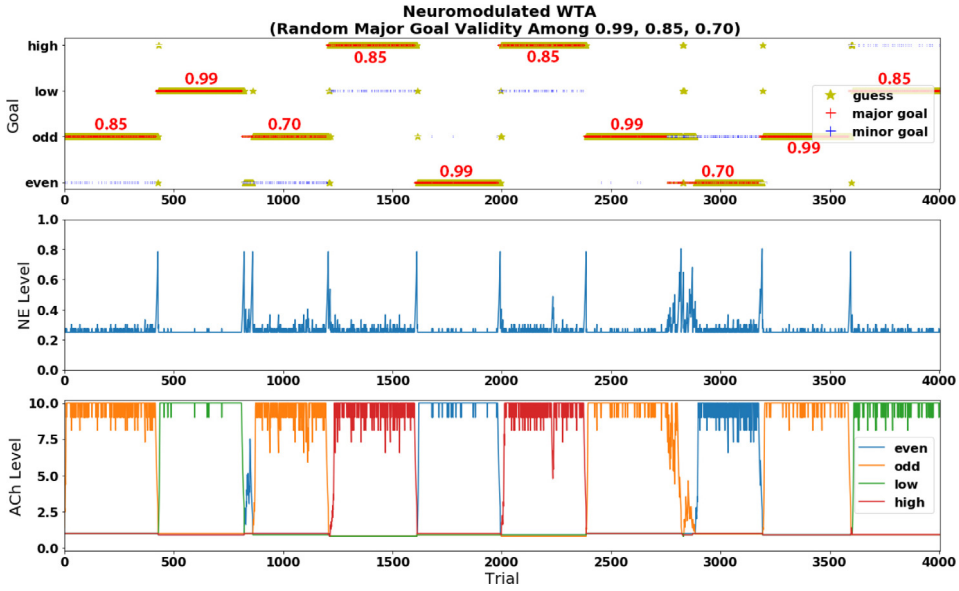


**Fig. 11.** Visualization of goal-driven perception performance on noisy MNIST pairs with the major goal validity randomly switching among 0.99, 0.85, and 0.70. In this neuromodulated benchmark, WTA replaced the softmax distribution in our model for cholinergic goal guessing. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

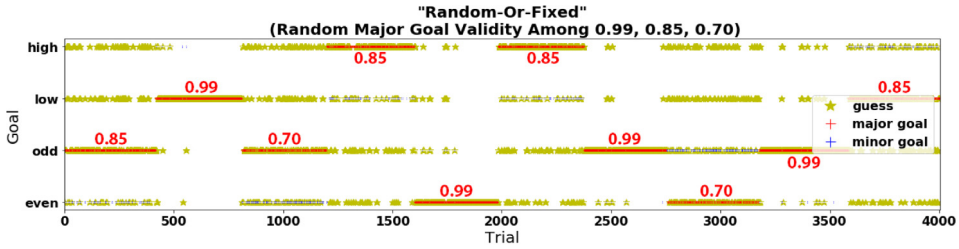


**Fig. 12.** Visualization of goal-driven perception performance on noisy MNIST pairs with the major goal validity randomly switching among 0.99, 0.85, and 0.70. In this "random-or-fixed" benchmark, whether the guessed goal was random or stayed the same depended on whether there was a mismatch or match in the goal guessing process of the previous trial. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

neuromodulated softmax and neuromodulated WTA, their overall accuracy of major and minor goal guessing was quite similar. However, comparing Fig. 9 for softmax with Fig. 11 for WTA, we observed that the softmax function allowed higher chances of selecting the minor goal during the intervals at which the major goal validity was low, whereas the WTA function would like to select the major goal regardless of its true validity and could cause a much longer lag when the major goal validity dropped.

### 3.5. Goal-driven perception on robot

To demonstrate that our model could generalize to a more practical application than MNIST digits, we tested our model on a human support robot that needed to guess an action with the neuromodulatory head, selectively attend to an object that corresponded to that action, and retrieve the object. The methods for this scenario were described in Section 2.3.

Fig. 13 shows the performance of a typical run of the action-based goal-driven perception. Its neuromodulation process was similar to Algorithm 1 for the noisy MNIST-pair experiment. Changes included using input images from three view angles of the HSR's camera, a fixed goal (action) validity of 1, and several parameter value adjustments (i.e., $\beta = 10$, $trial\_interval = 50$, $trial\_range = 0$, $ne_{correct} = 0.75$, $ne_{wrong} = 1.15$, $ch_{correct} = 1.35$, $ch_{wrong} = 0.95$). The average goal selection error for 5 runs was 23.8%, which was higher than the noisy MNIST-pair experiment (see Table 3) because of the shortened trial interval for each goal (action) switch. The average c-EB object prediction incorrectness was 30.6%. The average lag length was 13 trials. Uncertainties in each trial were addressed by possible object location switch, possible object removal and/or introduction, possible multi-instances of the same object(s), and slight view angle adjustment, in addition to possible true action switch (i.e., every 50 trials, not given to the agent). A complete trial with HSR in the testing room can be watched in a YouTube video (https://youtu.be/DUy-0fDZEvY).

## 4. Discussion

### 4.1. Main findings

In this paper, we have shown that a neuromodulated goal-driven perception model, which combines ideas from neuroscience with goal-driven perception in machine learning and artificial neural networks, could track context and flexibly shift attention to intended goals. Among many top-down attentional systems, we adapted c-EB (Zhang et al., 2018) as part of our model because of its similarities to how the ACh neuromodulatory system both increments attention to a goal and decrements attention to a distractor (Baxter & Chiba, 1999; Oros et al., 2014). Goals are often unknown and need to be discovered. The c-EB algorithm was modified to support multiple goals. After training, the biologically inspired algorithm could quickly learn the context without supervision, flexibly apply attention to the appropriate goal, and rapidly detect and re-adapt to context changes.

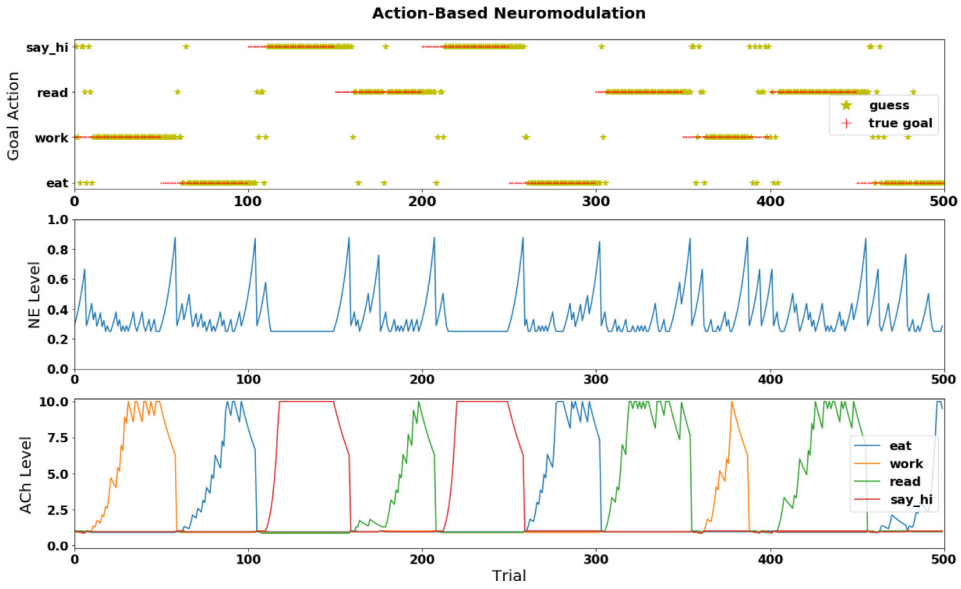


**Fig. 13.** Visualization of action-based goal-driven perception performance on different angles of robot views in an indoor scenario. The true goal action was randomly picked every 50 trials for 10 switches in a run. A softmax function (see Eq. (4), with $\beta = 10$) was applied to ACh levels for action guessing. See text for details. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

*4.1.1. Neural implementation of uncertainty tracking*

Yu and Dayan (2005) proposed a Bayesian model of neuromodulation in which the ACh system tracked expected uncertainty and the NE system tracked unexpected uncertainty. The present paper advances this work in two ways to support goal-driven perception: (1) The Bayesian model was recast as a neural model to make it compatible with neural networks. The neuromodulators were implemented as a neural network layer to drive attention toward a goal digit and divert attention away from distractors. (2) A neural network reset was implemented to rapidly re-adapt when a goal changes.

Neuroanatomical studies show the basal forebrain, which contains ACh neurons, has topographical connections specific to stimulus modalities and values (Zaborszky, 2002). Therefore, different ACh neurons tracked the expected uncertainties of different potential goals in our model. In a dynamic situation, the goal identity can change unexpectedly. Empirical evidence suggests that the NE system detects such changes and generates a "reset" signal to discard prior expectations when these expectations are violated (Bouret & Sara, 2005; Grella et al., 2019). In our experiments, the NE system rapidly recognized a change in the goal contingency, and drove a reset of ACh and NE activities. This caused the neural network to quickly explore new goals. It should be noted that the "reset" did not erase the learned object categories (e.g., digit parity and magnitude). Instead, it cleared the prior likelihood of potential goals, and resulted in a rapid re-adaptation to the new goal distribution.

In the real world, goals are often uncertain and unknown. In our noisy MNIST-pair experiment, the goal validity (i.e., probability of a goal being rewarded) ranged from 0.99 to 0.85 to 0.70, and the system needed to respond by either choosing the most likely goal or exploring alternative goals. Furthermore, the experimental design had a hierarchy of goals. For example, the goal would be to attend to the parity goal class and the sub-goal might be to reward odd digits 70% of the time and even digits 30% of the time. Interestingly, the neural network would often stay within a goal class (i.e., to choose parity and not magnitude).

In the robot action-based attention experiment, the objects linked with a predicted goal action might or might not exist in the views and may probably be at different locations. The adapted c-EB attention mechanism could pay significantly higher attention to existing objects. Selecting the highest attention region helped further with object localization and prediction. In both the MNIST-pair and real-scenario experiments, the unexpected major goal (action) switch after some trials could be quickly caught by the network within an acceptable lag.

*4.1.2. Exploration and uncertainty seeking*

Exploring options, rather than always choosing the most likely goal, is known as probability matching behavior (Wozny, Beierholm, & Shams, 2010). Similar to the results presented here, humans tend to underselect the most rewarding goal (Craig, Phillips, Zaldivar, Bhattacharyya, & Krichmar, 2016). Such behavior may be due to feature exploration, as subjects test hypotheses by switching between the features before deciding upon their most rewarding goal. In rodent studies, it has been shown that rats will seek uncertainty, and that this uncertainty seeking is governed by the ACh system (Naude et al., 2016). These uncertainty seeking strategies that appear in natural systems may be advantageous for artificial systems that are deployed in dynamic environments.

*4.2. Related work*

Top-down task-driven attention is an important mechanism for efficient visual search in humans and artificial systems (Baluch & Itti, 2011). Many computational models of attention have been proposed and implemented to either explain top-down attention or develop an application inspired by these mechanisms (Tanner & Itti, 2017, 2019; Tsotsos et al., 2015). Of particular interest are attentional systems that can leverage the power of CNNs. In these cases attentional information can propagate backwards, highlighting the features of a given goal (Zhang et al., 2018; Zhou, Khosla, Lapedriza, Oliva, & Torralba, 2016). Similar to the effect of the ACh system to increment attention to a goal and decrement attention to distractors (Baxter & Chiba, 1999), Zhang et al. (2018) proposed an Excitation Backprop (EB) mechanism with a contrastive top-down signal to enhance the perception of goal features. Similarly, Zhou et al. (2016) proposed a technique called Class Activation Mapping (CAM) for identifying regions in an attention map. Selvaraju et al. (2017) proposed Gradient-weighted

Class Activation Mapping (Grad-CAM) to highlight regions of interest and generate visual explanations. Their model could be applied to any CNN with no re-training. Similarly, our proposed model can work with any CNN. Moreover, our model replaces the Winner-Take-All mechanism or rigid probabilistic methods, with a flexible and adaptable layer based on neurobiological neuromodulation.

Intrinsic rewards and curiosity seeking have similarities to the exploration due to uncertainty demonstrated by our model. These intrinsic reward systems typically are rewarded for exploring infrequently observed states (Achiam & Sastry, 2017; Burda et al., 2018; Pathak, Agrawal, Efros, & Darrell, 2017), whereas the model introduced here selects goals based on the expected uncertainty of stimuli. In future work, it may be of interest to combine intrinsic rewards with uncertainty seeking.

#### 4.2.1. Uniqueness of our model

Our model is unique compared to existing attention models – they only focus on highlighting predefined (and pre-trained) goal objects in test images (Cao et al., 2015; Cho, Courville, & Bengio, 2015), without any ability to deal with unpredictable switching and validity of goals. c-EB, which we adapted in our work, has been shown to achieve top-down attention competitively and robustly (Zhang et al., 2018). Moreover, it is similar to how the ACh neuromodulatory system both increments attention to a goal and decrements attention to a distractor (Baxter & Chiba, 1999; Oros et al., 2014). The neuromodulatory layer on top of a top-down attentional network demonstrates a means toward goal-driven perception where the system can autonomously learn which objects to attend to and which objects to filter out in a noisy, dynamic setting.

In addition to the unique aspects of our neuromodulatory model, its robustness was ascertained via comparisons with neuromodulated WTA and "random-or-fixed" benchmarks. We also used ablation studies to show the necessity of having both ACh and NE neuromodulators to track the expected and unexpected uncertainties of goals and respond appropriately when the goal distribution changes. Further, generalization was ascertained via the HSR implementation in a real indoor scenario.

### 4.3. Future directions

#### 4.3.1. Handling new goals

In the present work, the goal classes were known, and the system guessed the appropriate goal given a goal identity and goal validity. However, the system might need to adapt to new goals or new goal classes. Adding multiple heads to the output layer of the network is one way to handle this. Instead of re-training the stimuli (e.g., digits or real objects), it would require some additional training for the new goal classes. However, the architecture might be more scalable with a single head that learns the goals online without any *a priori* assumptions. Similar to the present model, these unknown goals would initially be guessed. After sufficient reward feedback, the model would associate different goals with different reward likelihoods. The introduction of the ACh/NE neuromodulation should make the goal search fast and flexible. This will be explored in future iterations of our work.

#### 4.3.2. Different attentional mechanisms

The choice of c-EB for a top-down attentional mechanism was motivated by its similarity to ACh system and its affect on top-down attention. However, as mentioned above, we believe that the proposed system could also work with other state-of-the-art attentional mechanisms, including the CAM (Zhou et al., 2016) and its more general variation Grad-CAM (Selvaraju et al., 2017). As long as the neural network structure can support an additional neuromodulation layer, and there is some means to flow goal information from the top to lower layers, our neuromodulatory goal-driven perception system should be compatible.

#### 4.3.3. Application for artificial intelligence (AI)

We have shown the compatibility of the adapted c-EB attention mechanism with the Microsoft COCO dataset and an indoor scenario. Our model is applicable in broader AI scenarios. If a system (e.g., a self-driving car, a human support robot, etc.) faces many known and unknown task structures, our neuromodulatory goal-driven architecture would be able to help it choose tasks wisely regarding seen/unseen goals in a complex scenario.

#### 4.3.4. Inspiration for cognitive neuroscience

Our experimental design could be replicated in biological studies with non-human primates or rodents to investigate relevant neuromodulatory signals in the brain. We predict that NE neurons would increase phasic activity after a goal switch. Corbetta, Patel, and Shulman (2008) have shown that the locus coeruleus/norepinephrine system redirects attention from one object to another, and switches attention between networks. Attention is strongly modulated by acetylcholine through its projections to sensory cortex (Sarter, Hasselmo, Bruno, & Givens, 2005). Cholinergic activation has been shown to increase goal-driven attention in V1 by increasing the firing rate of neurons coding the attended objects (Goard & Dan, 2009; Herrero et al., 2008). It would be of interest to test whether ACh activity to V1 becomes somewhat random after phasic NE responses and if ACh modulation varies depending on goal validity.

The robot experiments highlight a somewhat unexplored aspect of attention. In addition to feature or spatial attention, attention is deployed to intended actions (for a review, see Atkinson, Simpson, & Cole, 2018). Recent results suggest that attention is required for both action planning and movement outcome monitoring (Mahon, Bendžiūtė, Hesse, & Hunt, 2018). In our robot experiments, an intended action led to attention to an object associated with the desired action. Such an attentional network could have benefits for human–robot interaction, especially when the intended actions could change due to context.

## 5. Conclusions

In this paper, we introduced a model of ACh and NE neuromodulation to perform goal-driven perception. The proposed network architecture discovers goals using online learning, and highlights the stimulus features corresponding to the goal. Moreover, the proposed system rapidly adapts when goal contingencies change. This neurobiologically inspired model can be applied to other problem domains and other top-down attentional networks.

## Acknowledgments

This material is based upon work supported by the United States Air Force and DARPA under Contract No. FA8750-18-C-0103, and other support in part by Toyota Motor North America. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force and DARPA.

## References

Achiam, J., & Sastry, S. (2017). Surprise-based intrinsic motivation for deep reinforcement learning. arXiv preprint arXiv:1703.01732.

Atkinson, M. A., Simpson, A. A., & Cole, G. G. (2018). Visual attention and action: How cueing, direct mapping, and social interactions drive orienting. *Psychonomic Bulletin & Review*, *25*(5), 1585–1605.

Avery, M. C., Dutt, N., & Krichmar, J. L. (2014). Mechanisms underlying the basal forebrain enhancement of top-down and bottom-up attention. *European Journal of Neuroscience*, *39*(5), 852–865. http://dx.doi.org/10.1111/ejn.12433.

Baluch, F., & Itti, L. (2011). Mechanisms of top-down attention. *Trends in Neurosciences*, *34*(4), 210–224.

Baxter, M. G., & Chiba, A. A. (1999). Cognitive functions of the basal forebrain. *Current Opinion in Neurobiology*, *9*(2), 178–183.

Bouret, S., & Sara, S. J. (2005). Network reset: A simplified overarching theory of locus coeruleus noradrenaline function. *Trends in Neurosciences*, *28*(11), 574–582.

Burda, Y., Edwards, H., Pathak, D., Storkey, A., Darrell, T., & Efros, A. A. (2018). Large-scale study of curiosity-driven learning. arXiv preprint arXiv:1808.04355.

Cao, C., Liu, X., Yang, Y., Yu, Y., Wang, J., Wang, Z., et al. (2015). Look and think twice: Capturing top-down visual attention with feedback convolutional neural networks. In *Proceedings of the IEEE international conference on computer vision* (pp. 2956–2964).

Cho, K., Courville, A., & Bengio, Y. (2015). Describing multimedia content using attention-based encoder-decoder networks. *IEEE Transactions on Multimedia*, *17*(11), 1875–1886.

Corbetta, M., Patel, G., & Shulman, G. L. (2008). The reorienting system of the human brain: From environment to theory of mind. *Neuron*, *58*(3), 306–324.

Craig, A. B., Phillips, M. E., Zaldivar, A., Bhattacharyya, R., & Krichmar, J. L. (2016). Investigation of biases and compensatory strategies using a probabilistic variant of the Wisconsin Card Sorting Test. *Frontiers in Psychology*, *7*(17), http://dx.doi.org/10.3389/fpsyg.2016.00017.

Goard, M., & Dan, Y. (2009). Basal forebrain activation enhances cortical coding of natural scenes. *Nature Neuroscience*, *12*(11), 1444.

Grella, S. L., Neil, J. M., Edison, H. T., Strong, V. D., Odintsova, I. V., Walling, S. G., et al. (2019). Locus coeruleus phasic, but not tonic, activation initiates global remapping in a familiar environment. *Journal of Neuroscience*, *39*(3), 445–455.

Greydanus, S. (2018). *Excitationbp: Visualizing how deep networks make decisions*. GitHub. MIT License, https://github.com/greydanus/excitationbp.

Herrero, J. L., Roberts, M., Delicato, L. S., Gieselmann, M. A., Dayan, P., & Thiele, A. (2008). Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature*, *454*(7208), 1110.

Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, *40*(10–12), 1489–1506.

Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., et al. (2014). Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM international conference on multimedia* (pp. 675–678). ACM.

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, *86*(11), 2278–2324.

Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., et al. (2014). Microsoft coco: Common objects in context. In *European conference on computer vision* (pp. 740–755). Springer.

Mahon, A., Bendžiūtė, S., Hesse, C., & Hunt, A. R. (2018). Shared attention for action selection and action monitoring in goal-directed reaching. *Psychological Research*, 1–14.

Nair, V., & Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning* (pp. 807–814).

Naude, J., Tolu, S., Dongelmans, M., Torquet, N., Valverde, S., Rodriguez, G., et al. (2016). Nicotinic receptors in the ventral tegmental area promote uncertainty-seeking. *Nature Neuroscience*, *19*(3), 471–478. http://dx.doi.org/10.1038/nn.4223.

Oros, N., Chiba, A. A., Nitz, D. A., & Krichmar, J. L. (2014). Learning to ignore: A modeling study of a decremental cholinergic pathway and its influence on attention and learning. *Learning & Memory*, *21*(2), 105–118.

Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., et al. (2017). Automatic differentiation in PyTorch. In *NIPS autodiff workshop*.

Pathak, D., Agrawal, P., Efros, A. A., & Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops* (pp. 16–17).

Sarter, M., Hasselmo, M. E., Bruno, J. P., & Givens, B. (2005). Unraveling the attentional functions of cortical cholinergic inputs: Interactions between signal-driven and cognitive modulation of signal detection. *Brain Research Reviews*, *48*(1), 98–111.

Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *2017 IEEE international conference on computer vision* (pp. 618–626). IEEE.

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., et al. (2015). Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1–9).

Tanner, J., & Itti, L. (2017). Goal relevance as a quantitative model of human task relevance. *Psychological Review*, *124*(2), 168.

Tanner, J., & Itti, L. (2019). A top-down saliency model with goal relevance. *Journal of Vision*, *19*(1), 1–16.

Tsotsos, J. K., Eckstein, M. P., & Landy, M. S. (2015). Computational models of visual attention. *Vision Research*, *116*(Pt B), 93–94. http://dx.doi.org/10.1016/j.visres.2015.09.007.

Wozny, D. R., Beierholm, U. R., & Shams, L. (2010). Probability matching as a computational strategy used in perception. *PLoS Computational Biology*, *6*(8), http://dx.doi.org/10.1371/journal.pcbi.1000871.

Yamamoto, T., Nishino, T., Kajima, H., Ohta, M., & Ikeda, K. (2018). Human support robot (HSR). In *ACM SIGGRAPH 2018 emerging technologies* (p. 11). ACM.

Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, *46*(4), 681–692. http://dx.doi.org/10.1016/j.neuron.2005.04.026.

Zaborszky, L. (2002). The modular organization of brain systems. Basal forebrain: The last frontier. In *Progress in brain research (vol. 136)* (pp. 359–372). Elsevier.

Zhang, J., Bargal, S. A., Lin, Z., Brandt, J., Shen, X., & Sclaroff, S. (2018). Top-down neural attention by excitation backprop. *International Journal of Computer Vision*, *126*(10), 1084–1102.

Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., & Torralba, A. (2016). Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2921–2929).