NeuroGym: An open resource for developing and sharing neuroscience tasks

Manuel Molano-Mazón[1,2], Joao Barbosa[2], Jordi Pastor-Ciurana[1], Marta Fradera[1], Ru-Yuan Zhang[3], Jeremy Forest[4], Jorge del Pozo[1], Li Ji-An[5], Christopher J Cueva[6], Jaime de la Rocha[1], Devika Narain[7], Guangyu Robert Yang[8]

[1] Brains Circuits and Behavior lab, IDIBAPS, Spain
[2] Group for Neural Theory, Ecole Normale Supérieure, Paris, France
[3] Institute of Psychology and Behavioral Science, Shanghai Jiao Tong University
[4] Cornell University
[5] University of California San Diego
[6] Columbia University, United States
[7] Dept. of Neuroscience, Erasmus University Medical Center, Rotterdam, The Netherlands
[8] Department of Brain and Cognitive Sciences, MIT

Corresponding authors; email: manuelmolanomazon@gmail.com, yanggr@mit.edu

Abstract

Artificial Neural Networks (ANNs) trained on specific cognitive tasks have re-emerged as a useful tool to study the brain. However, ANNs would better aid cognitive neuroscience if a given network could be easily trained on a wide range of tasks for which neural recordings are available. Moreover, unintentional divergent implementations of cognitive tasks can produce variable results, which limits their interpretability. Towards this goal, we present NeuroGym, an open-source Python package that provides a large collection of customizable neuroscience tasks to test and compare network models. Building upon the OpenAI Gym toolbox, NeuroGym tasks (1) are written in a high-level flexible Python framework; (2) possess a shared interface tailored to common needs of neuroscience tasks that facilitates their design and usage; (3) support the training of ANNs using both Reinforcement and Supervised Learning techniques. The toolbox allows easy assembly of new tasks by modifying existing ones in a hierarchical and modular fashion. These design features make it straightforward to take a network designed for one task and train it on many other tasks. NeuroGym is a community-driven effort that contributes to a rapidly evolving open ecosystem of neural network development, data analysis, and model-data comparison.

Introduction

In the last decade, the renewed use of Artificial Neural Networks (ANNs) in computational neuroscience have allowed modeling the behavior of animals performing complex tasks, helping to generate hypotheses about the neural mechanisms underlying such behavior (Mante et al. 2013; Sussillo et al. 2015; Carnevale et al. 2015; Chaisangmongkon et al. 2017; Cueva and Wei 2018; J. Wang et al. 2018; Remington et al. 2018; Mastrogiuseppe and Ostojic 2018; Farrell et al. 2019; Masis et al. 2020; Saxena et al. 2021; Feulner and Clopath 2021; Schaeffer et al., 2020). In turn, experimental advances such as large scale

recordings obtained from animals performing complex behavioral tasks, produce ever growing, multi-area and multi-cell-type datasets that demand a continuous reevaluation on the ANNs used to model them. However, a key requirement to fully exploit the potential of the interaction between experiments and modeling is to develop engineering tools that allow for an easy comparison of the results obtained across different neural systems and experimental paradigms (Fig. 1).

The use of ANNs in cognitive neuroscience typically consists of training a network in a small set of closely related tasks. However, a major goal when training neural networks must be to find a model that can explain a wide range of experimental results collected across many different tasks. A necessary step is to put together a large set of neuroscience tasks on which different models can be trained. Indeed, there is a large body of experimental work that hinges on a number of canonical behavioral tasks and their variants, such as perceptual decision making and working memory tasks. This accumulated work makes it possible to develop a common computational framework encompassing many relevant tasks on which neural networks can be trained.

Here, we introduce an easy-to-use toolkit called NeuroGym, that allows training any network model on many established neuroscience tasks. NeuroGym builds upon OpenAI Gym (Brockman et al. 2016), the quintessential collection of reinforcement learning environments in Machine Learning. NeuroGym introduces several new features: **1)** a high-level task constructor that allows users to more easily build their own neuroscience tasks and adapt existing ones; **2)** a variety of tools called wrappers that permit assembling complex tasks using simpler ones as modules; **3)** support for training models using either Reinforcement (RL) or Supervised Learning (SL) **4)** programs purely based on Python, which makes it easy to use and code.

Here we provide a developer-oriented summary of the core structure of NeuroGym and of each of its main components. Our aim is to periodically update this manuscript with the hope that it will serve as a guide to use and create NeuroGym tasks, and will help to create, test and compare network models of biological neural circuits.
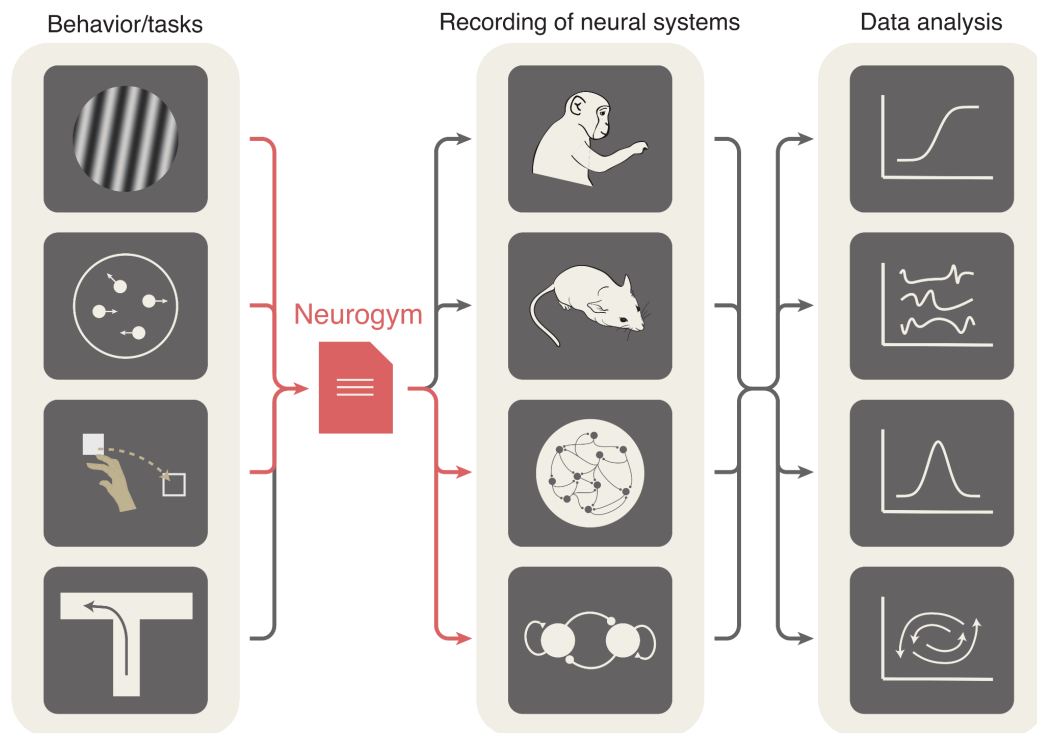
Figure 1. To fully exploit the new advances in experimental neuroscience, a necessary step is to develop an open ecosystem in which many neural systems, both biological and *in silico* (middle column), can be trained in many standardized behavioral paradigms (left column) and be analyzed with the techniques commonly used in computational neuroscience (right column). NeuroGym aims at contributing to such an ecosystem by providing a curated collection of neuroscience tasks (see Appendix I) that supports easy training of a given model in many cognitive tasks (pink arrows), hence speeding up the process of developing neural network models that can be compared to large-scale recordings and behavioral datasets.

How it works

Tasks structure

Neuroscience tasks vary widely at the computational level, which reflects the diversity of cognitive functions that they test. However, there are several commonalities among them. This permits us to define a core structure shared by all tasks, greatly facilitating their usage and the construction of new ones.

Ultimately, all experimental neuroscience tasks are designed to investigate the behavior of a biological neural system in an environment defined by a set of parameters. These parameters can typically vary at three different timescales. First, some parameters will remain constant throughout the entire experiment, defining the general properties of the environment. Other parameters will vary periodically in intervals called *trials*. The organization of an experiment into trials facilitates the interpretation of the responses of a neural system in a given environment by allowing the exploration of different parameter configurations. The rest of parameters will vary within the trials to assure variability among trials. For instance in the Random Dots Motion (RDM) task (Britten et al. 1992), subjects have to report the direction in which a cloud of dots is moving (Fig. 2). In this case, the parameters involved are: 1) the ranges within which the dots positions and velocities will vary. These ranges usually remain constant during the entire task; 2) the distribution of velocities (including direction and magnitude) for the dots at every trial. It is this distribution of velocities that the subjects has to infer and respond to, and it is the responses they provide that are the object of the experiment; 3) The exact positions of each dot at each frame composing a single trial are decided based on the distribution defined for each trial and will be randomly selected for each trial.

We will now build a NeuroGym task to illustrate its basic structure. In NeuroGym, tasks are built as python classes composed of three main functions that run at the three timescales described above: ___init___, _*new_trial*_ and _*step*_.

Defining the general properties of the task (__init__ function)

The first step to build a task that tests a specific cognitive aspect of a neural system is to define two attributes: the type and range of the observations provided to the system and the actions with which the system will respond. In NeuroGym, these attributes are defined in the ___init___ function using the *Space* toolbox (Brockman et al. 2016) (Fig. 3). For instance in the example code shown in Figure 3 (lines 10-14), the observations (*observation_space*) will be three-dimensional with unbounded float values, while the actions (*action_space*) will be discrete, taking one of three possible values.
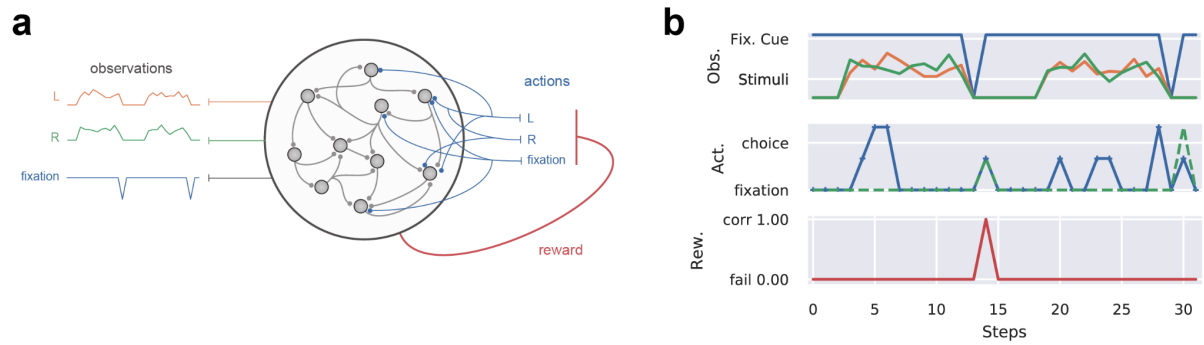
Figure 2. Two consecutive trials of the RDM task. Here, the velocities towards the left and right sides have been abstracted as two noisy signals with different means. The model has to choose the action (Left/Right) associated with the signal that is largest on average. At each timestep the task provides an observation consisting of two noisy stimuli (Fig. 2, orange and green traces) and a fixation cue (Fig. 2, gray thick) (top panel), and receives an action from the model (middle panel, blue, solid trace). In response to each action the task also provides a scalar value (the reward (bottom panel, red trace)) that serves as feedback to the model. The task can also provide additional information, useful for posterior analyses (e.g. the ground-truth at each timestep, (middle panel, green trace)). In this example, the model corresponds to an LSTM (Hochreiter and Schmidhuber 1997) network trained using Reinforcement Learning (A2C algorithm (Mnih et al. 2016)).

The **_init_** function (Fig. 3) is run at the moment of creating the task and, apart from the observations and actions spaces, it defines any other general property of the environment in which the subject will behave. One important feature that is present in many neuroscience tasks is the time periods. For instance, in the RDM task there are three different periods: 1) an initial fixation period during which the subject typically has to fixate the gaze on the center of the screen; 2) a stimulation period during which the moving dots are presented; 3) a decision period in which the subject has to report is choice. The duration and ranges of these periods are defined in the __init__ function (*timing*, line 7). As it is typically done, NeuroGym simulates real time with discrete timesteps whose length can be input to the __init__ function as *dt.* Therefore, for each trial, each period will be composed of a number of timesteps that may have a constant value throughout the experiment, or may vary from trial to trial (e.g. a variable stimulation period as in line 8), taking values drawn from a distribution also defined in the **_init_** function and rounded to a multiple of *dt*.

```
1   class RDM(ngym.TrialEnv):
2       def __init__(self, dt=100):
3           super().__init__(dt=dt)
4           self.cohs = [0, 6.4, 12.8, 25.6, 51.2]   # Coherences values
5           self.rewards = {'abort': -0.1, 'correct': +1., 'fail': 0.}   # Rewards
6           # Periods duration
7           self.timing = {'fixation': 500,
8                          'stimulus': ngym.random.TruncExp(330, 80, 1500, rng=self.rng),
9                          'decision': 500}
10          # Action and observation spaces
11          name = {'fixation': 0, 'stimulus': range(1, 3)}
12          self.observation_space = spaces.Box(-np.inf, np.inf, shape=(3,), dtype=np.float32, name=name)
13          name = {'fixation': 0, 'choice': range(1, 3)}
14          self.action_space = spaces.Discrete(3, name=name)
```

Figure 3. Implementation of the Random Dots Motion task (Fig. 2). A noisy stimulus is shown during the stimulation period whose strength (coherence) is randomly sampled every trial. Because the stimulus is noisy, the agent has to integrate the stimulus over time. **_init_** **function**: Observation and action ranges are defined using the *Space* toolbox (Brockman et al. 2016): (1) *observation_space* describes the format and range of the observations that will be provided to the model (i.e. the model inputs); (2) *action_space* describes the format and ranges of valid actions (i.e. the model outputs). NeuroGym supports setting the time length of each step (*self.dt*, in ms), and the time length of each time period (*self.timing*).

Task interface (_step function)

Once the main structure of the task is defined, we may want to specify how the observations and rewards provided by the task depend on the actions made by the model. For instance, in the RDM task, if the model chooses left when the dots are moving to the right, the task may provide a punishment in the next timestep (Fig. 4, lines 41-47). This corresponds to defining the rules that the model must learn and follow to increase the amount of reward it receives. These rules are defined in the **_step** function, which receives an action at each timestep (Fig. 2, middle panel, blue trace; Fig. 4 line 32) and outputs a new observation (Fig. 2, top panel) and a reward (Fig. 2, bottom panel, red trace). Additionally, it outputs two variables that contain meta-information useful for debugging and offline analysis: a boolean variable describing whether the environment needs to be reset (*done*), and a dictionary holding any additional information (*info*) (Fig. 4, line 48).

```
32    def _step(self, action, **kwargs):
33          new_trial = False
34          obs = self.ob_now
35          gt = self.gt_now
36          reward = 0
37          if self.in_period('fixation'):  # If in fixation period
38              if action != 0:
39                  reward = self.rewards['abort']
40          elif self.in_period('decision'):  # If in decision period
41              if action != 0:
42                  new_trial = True
43                  if action == gt:
44                      reward = self.rewards['correct']
45                      self.performance = 1
46                  else:
47                      reward = self.rewards['fail']
48          return obs, reward, False, {'new_trial': new_trial, 'gt': gt}
```

Figure 4. _step function: During *_step*, the entire observations and ground truth can be accessed via the *self.ob* and *self.gt* attributes, and their current values with *self.ob_now* and *self.gt_now*. The *self.in_period* method allows checking the current period.

Although the **_init_** and **_step** functions fully define a task, most neuroscience tasks are organized into trials. In NeuroGym this is done using the **_new_trial** function, which we explain next.

Trial-structured tasks (_new_trial function)

To facilitate the creation of trial-structured tasks, NeuroGym includes the **_new_trial** function. The **_new_trial** function generates all the necessary information (e.g. stimulus coherence in the RDM task) to generate the observations and rewards that will be provided to the model during the trial. It takes any key-word arguments that can be used to modify the task variables that were defined in **__init__** (see below), and creates a dictionary trial that contains all relevant information about the new incoming trial and can be accessed through the entire trial.

Depending on the algorithm used to train the model, the _new_trial function can serve different purposes. In reinforcement-learning setups (Sutton and Barto 2018), the model communicates with the task exclusively via the **_step** function, choosing actions and receiving observations and rewards that depend on those actions. In this scenario, the **_new_trial** function will be run at the beginning of each trial. However, artificial neural networks are often trained using supervised learning (Werbos 1990), which is typically much faster than reinforcement learning, at the cost of parting from the way biological circuits learn. In this case, it is actually possible to construct the dataset with observations and their associated labels *a priori*, bypassing the **_step** function and only using **_new_trial**. Note however that this approach is not always possible, as many tasks will not allow pre-defined observations or the associated rewards. This is the case whenever observations or rewards depend on the actions made by the model, such as tasks measuring reaction times (Roitman and Shadlen 2002).

```
15    def _new_trial(self, **kwargs):
16        # Set rewarded side (ground truth)
17        ground_truth = self.rng.choice(np.arange(2))
18        trial = {'ground_truth': ground_truth, 'coh': self.rng.choice(self.cohs)}
19        # Set periods
20        self.add_period(period='fixation')
21        self.add_period(['stimulus', 'decision'], after='fixation')
22        # Set stimuli
23        self.add_ob(1, 'fixation', where='fixation')
24        stim = np.ones(2) * (1 - trial['coh']/100)/2
25        stim[trial['ground_truth']] = (1 + trial['coh']/100)/2
26        self.add_ob(stim, 'stimulus', where='stimulus')
27        #  Add noise to stimulus observations
28        self.add_randn(0, 1/np.sqrt(self.dt), 'stimulus', where='stimulus')
29        # Set ground-truth vector
30        self.set_groundtruth(ground_truth, period='decision', where='choice')
31        return trial
```

Figure 5. **_new_trial** function: creates a dictionary that contains all relevant information about the new incoming trial and can be accessed as *self.trial*. Observation and ground-truth are set with the self.*add_ob* and *self.set_groundtruth* methods. The periods are added through the *self.add_period* method.

Modifying existing tasks (wrappers)

Most neuroscientific projects involve the use of several variants of the same task. Although this can be often achieved by directly adjusting the parameters that define the task (e.g. the stimulation period duration), other variants require modifying its logic. For instance, the RDM task described above could be transformed into a reaction-time task by allowing the model to end the trial at any time during the stimulation period.

There are several approaches to create variants of the same task: (1) Implement two separate tasks from scratch that only differ in a few details. (2) Use conditional code to deal with the two variants separately (e.g. an **if reaction_time then**). (3) Construct the tasks in a modular fashion, implementing the parts that are common to all variants in a core task and invoking it from the different variants. While the first two options are feasible for most neuroscientific projects that revolve around a single paradigm, scaling up to hundreds of tasks is only feasible with the third approach. In NeuroGym this is achieved by leveraging the functionality called wrapper (Brockman et al. 2016), which is a Python class that allows modifying existing tasks in potentially infinite ways (see Appendix II).

The simplest way of modifying an existing task using wrappers is to only change the input or the output of the **_step** function. In this case, wrappers can inherit directly from the **gym.Wrapper** class present in the Gym toolbox. For instance, one can pass the previous action and previous reward to the model as part of the observation, which has been shown to allow the neural network models to learn at two different timescales, in what is called meta reinforcement learning (J. X. Wang et al. 2018).

```python
1  import gym
2  import neurogym.wrappers as wrappers
3  def roitman02(**kwargs):
4      env = gym.make('PerceptualDecisionMaking-v0', **kwargs)
5      env = wrappers.ReactionTime(env)
6      env = wrappers.PassReward(env)
7      return env
```

Figure 6. Implementation of a reaction-time task (Roitman and Shadlen 2002).

Other modifications require changing some of the latent variables that are used to build a new trial in the **_new_trial** method, such as the probability of the reward being associated to Left and Right sides, effectively introducing serial correlations in the task (Abbott et al. 2017; Molano-Mazon et al., 2021). To address these cases, NeuroGym extends the **gym.Wrapper** functionality with a new class, **neurogym.TrialWrapper**, that allows passing input arguments to the **new_trial** function of the task. Note that Trial-wrappers allow introducing variations in the task at an extra timescale, between the constant variables set in **__init__** and the trial-to-trial modifications done by **_new_trial**, effectively organizing the task into blocks.

Importantly, thanks to the modular approach of NeuroGym, wrappers can be easily combined in a few lines of code (see code in Fig. 4) greatly expanding the space of available tasks.
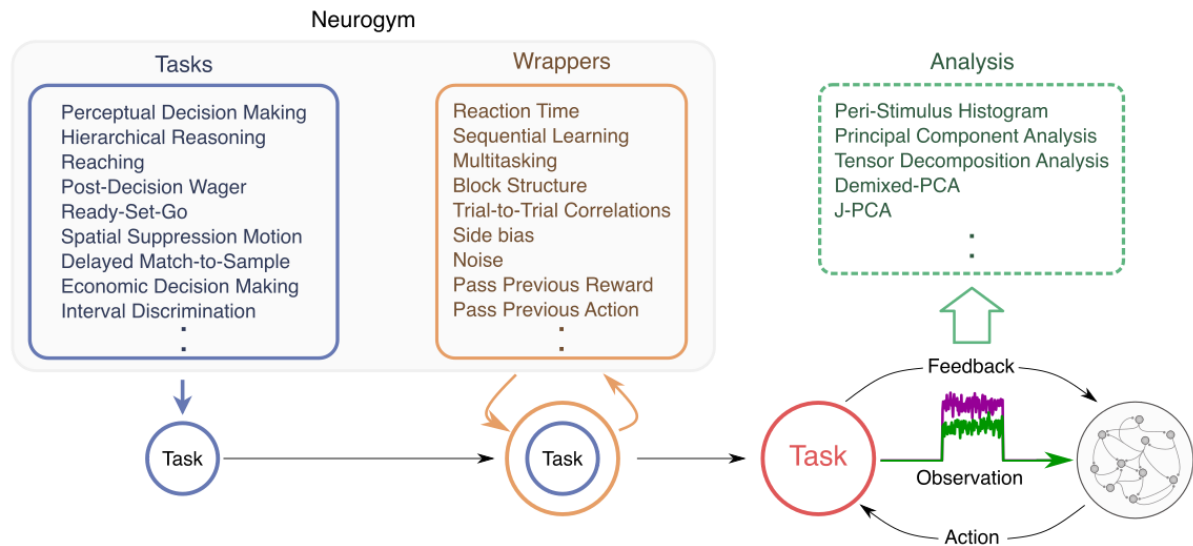
Figure 7. Schematic illustrating the NeuroGym toolbox pipeline. A task can be chosen from the more than 30 cognitive, sensory, and motor tasks already included in NeuroGym (purple box). The tasks can be modified using any number of wrappers (orange box). The model interacts with the final task (pink) by providing an action and receiving an observation and feedback, if applicable. The behavioral and neural data obtained from the training and testing can then be compared to experimental data using additional analysis code we provide in a separate repository (green box).

Applications

In this section, we illustrate the usage of NeuroGym with two real-world applications. Specifically, we trained RNN to perform two tasks that are not in the NeuroGym core ecosystem and compare their neural dynamics with electrophysiological recordings from animals performing similar tasks. In the first example, we analyze single-unit recordings from the prefrontal cortex of two macaques performing a working memory task (*oculomotor delayed-response task*, Figure 8a, (Barbosa et al. 2020)). We then build a new NeuroGym task simulating the one performed by the monkeys and train an RNN to do it using supervised learning. Finally, we compare the neural dynamics recorded from the monkey PFC to those in the RNN units . Then, we train a network in an abstraction of the task used by the International Brain Lab. We compare the network dynamics of the network with neuropixel recordings from mice performing a two-alternative, go-nogo forced choice task.

Oculomotor delayed-response task (ODR)

In the original task, monkeys had to fixate on the center of a screen while performing the task. During stimulus, a disk could appear at 1 of 8 possible locations (n_locations), after which there was a delay of 3 seconds long in which fixation had to be maintained. During the decision, the monkey trained to report the stimulus location (Fig. 8a).

To simulate the ODR task with NeuroGym, only the __init__ and _new_trial functions need to be defined, because the networks will be trained using supervised learning (see also *Task structure*). The task can be designed with a 2-dimensional observation_space and 2-dimensional action_space (representing the x,y location of the stimulus/report, respectively). Since the original study focused on what happened in between trials, we include a large inter-trial interval (ITI, 2000 timesteps). Next, we adapt _new_trial to reflect the structure of the task. Note that we are essentially simulating the monkey gaze, with an output of (0,0) for fixating at the center of the screen. Analogously, the output should have the same x,y coordinates as the stimulus during the decision period. The code to define the task and train a network can be found in our github. In Fig 8 we show the monkey reports and normalized firing rate of their PFC as well as from 64 units of the trained network.
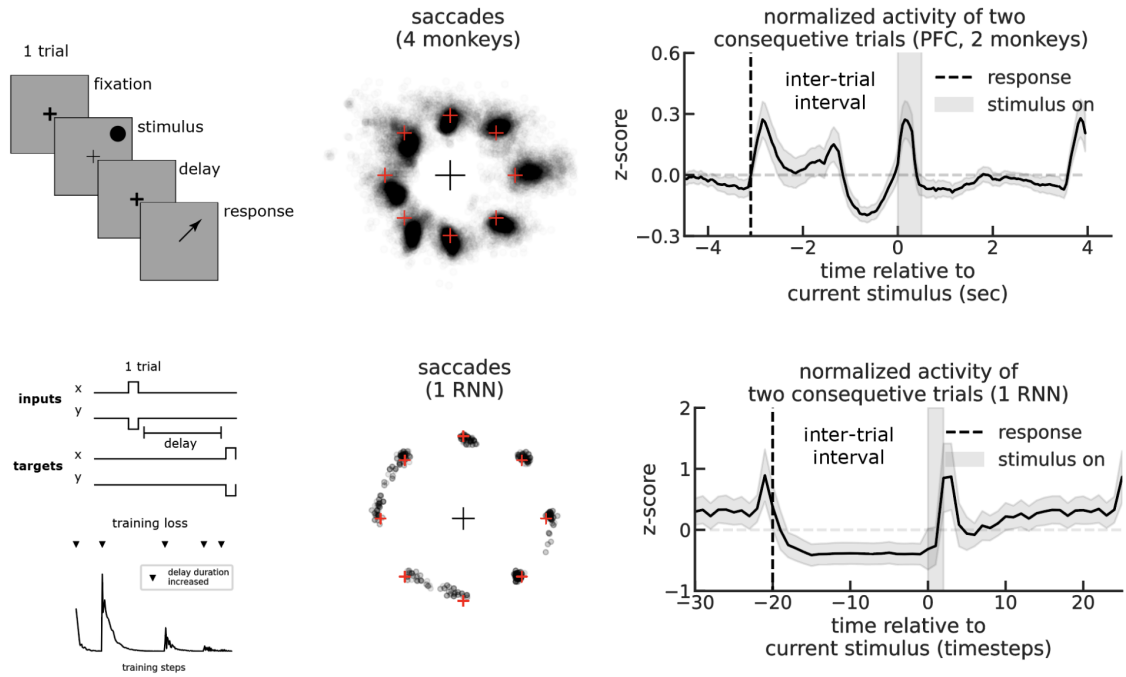
Figure 8. Top. Left, General task design. **Middle,** The distribution of eye positions following the saccade in the visually guided saccade task **Right**, Averaged, normalized firing rate of n = 206 neurons during the ITI (spike counts of 300-ms causal square kernel, z-scored in the interval [–4.5 s, 1.5 s]). Gray vertical bars mark the response and stimulus cue periods. **Bottom.** Same for the training of a single RNN in the NeuroGym task.

Two-alternative, go-nogo forced choice task

The development of biologically relevant models will need to be paralleled by the assembly of a common language that describes the neural computations used by biological and *in silico* neural systems. Towards such a common language, a key step will be the development of a set of core analysis methods organized into high-throughput pipelines that allow for rapid comparison of neural systems. With this goal in mind, we open sourced another codebase, called **ngym-usage**, for efficiently comparing networks with data through a variety of analysis methods. To illustrate the potential of this approach we trained RNNs to do a two-alternative, go-nogo forced choice task and compared the activity of the networks with the one recorded from mice performing a similar task (Steinmetz et al. 2019).

In the original task, mice had to turn a wheel with their forepaws to report which of two visual gratings presented in different computer screens had higher contrast. The trial structure was as follows: after mice held the wheel still for a short variable interval (0.2-0.5s), visual stimuli were presented. After stimulus onset there was a random delay interval (0.5-1.2s), during which mice could turn the wheel, but visual stimuli were locked in place. At the end of the delay interval, an auditory tone cue was delivered after which mice could report their choice using the wheel.

We simulated the task with NeuroGym by representing the gratings as two noisy signals drawn from a normal distribution with different means. Therefore, as above, the task can be

designed with a 2-dimensional observation_space and 3-dimensional action_space (representing fixation and turning left and right). We trained a 64-unit vanilla RNN on the task using supervised learning with Pytorch (Paszke et al. 2019) and compared the activity of the units in the network and that obtained from the recordings done by Steinmetz et al using dPCA (Kobak et al. 2016).
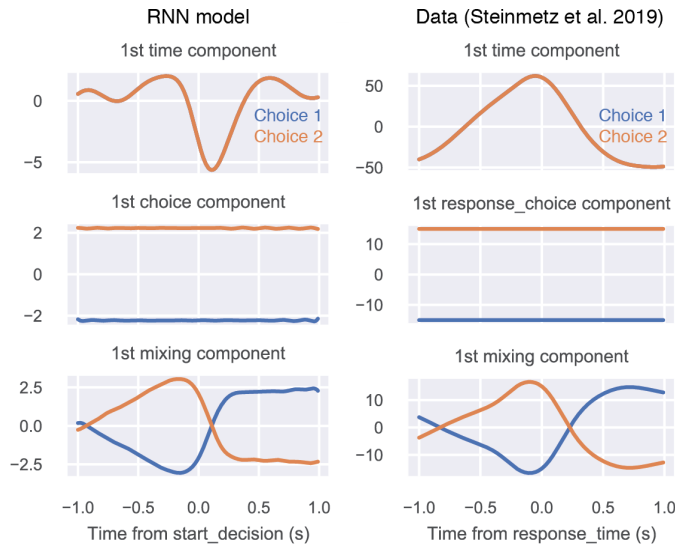

Figure 9. Demixed PCA (dPCA) (Kobak et al. 2016) analysis on both a network trained with SL and experimental mouse data on similar tasks (Steinmetz et al. 2019). The same analysis can be easily applied to both model and data since "neural recordings" are stored in the NWB format.

Acknowledgements

Appendix I. Core tasks

Anti-response task (Munoz and Everling 2004)
During the fixation period, the agent fixates on a fixation point. During the following stimulus period, the agent is then shown a stimulus away from the fixation point. Finally, the agent needs to respond in the opposite direction of the stimulus during the decision period.

Multi-arm bandit task (J. X. Wang et al. 2018)
On each trial, the agent is presented with multiple choices. Each option produces a reward of a certain magnitude given a certain probability.

Context-dependent decision-making task (Mante et al. 2013)
The agent simultaneously receives stimulus inputs from two modalities (for example, a colored random dot motion pattern with color and motion modalities), and has to make a perceptual decision based on only one of the two modalities, while ignoring the other. The agent reports its decision during the decision period, with an optional delay period in between the stimulus period and the decision period. The relevant modality is not explicitly signaled.

Two-step task (Daw et al. 2011)
On each trial, an initial choice between two options lead to either of two, second-stage states. In turn, these both demand another two-option choice, each of which is associated with a different chance of receiving reward.

Delay comparison task (Barak, Tsodyks, and Romo 2010)
The agent has to compare the magnitude of two stimuli separated by a delay period. The agent reports its decision of the stronger stimulus during the decision period.

Delayed match-to-category task (Freedman and Assad 2006)
A sample stimulus is shown during the sample period. The stimulus is characterized by a one-dimensional variable, such as its orientation between 0 and 360 degree. This one-dimensional variable is separated into two categories (for example, 0-180 degree and 180-360 degree). After a delay period, a test stimulus is shown. The agent needs to determine whether the sample and the test stimuli belong to the same category, and report that decision during the decision period.

Delayed match-to-sample task (Miller, Erickson, and Desimone 1996)
A sample stimulus is shown during the sample period. The stimulus is characterized by a one-dimensional variable, such as its orientation between 0 and 360 degree. After a delay period, a test stimulus is shown. The agent needs to determine whether the sample and the test stimuli are equal, and report that decision during the decision period.

Delayed match-to-sample task (with multiple, potentially repeating distractors) (Miller, Erickson, and Desimone 1996)
A sample stimulus is shown during the sample period. The stimulus is characterized by a one-dimensional variable, such as its orientation between 0 and 360 degree. After a delay period, the first test stimulus is shown. The agent needs to determine whether the sample and this test stimuli are equal. If so, it needs to produce the match response. If the first test

is not equal to the sample stimulus, another delay period and then a second test stimulus follow, and so on.

Delayed paired-association task (Zhang et al. 2019)
The agent is shown a pair of two stimuli separated by a delay period. For half of the stimuli-pairs shown, the agent should choose the Go response. The agent is rewarded if it chose the Go response correctly.

Two-item delay-match-to-sample task (Rose et al. 2016)
The trial starts with a fixation period. Then during the sample period, two sample stimuli are shown simultaneously. Followed by the first delay period, a cue is shown, indicating which sample stimulus will be tested. Then the first test stimulus is shown and the agent needs to report whether this test stimulus matches the cued sample stimulus. Then another delay and then test period follows, and the agent needs to report whether the other sample stimulus matches the second test stimulus.

Economic decision making task (Padoa-Schioppa and Assad 2006)
An agent chooses between two options. Each option offers a certain amount of juice. Its amount is indicated by the stimulus. The two options offer different types of juice, and the agent prefers one over another.

Go/No-go task (Zhang et al. 2019)
A stimulus is shown during the stimulus period. The stimulus period is followed by a delay period, and then a decision period. If the stimulus is a Go stimulus, then the subject should choose the action Go during the decision period, otherwise, the subject should remain fixating.

Hierarchical reasoning task (Sarafyazd and Jazayeri 2019)
On each trial, the subject receives two flashes separated by a delay period. The subject needs to judge whether the duration of this delay period is shorter than a threshold. Both flashes appear at the same location on each trial. For one trial type, the network should report its decision by going to the location of the flashes if the delay is shorter than the threshold. In another trial type, the network should go to the opposite direction of the flashes if the delay is short. The two types of trials are alternated across blocks, and the block transition is unannounced.

Interval discrimination task (Genovesio, Tsujimoto, and Wise 2009)
Comparing the time length of two stimuli. Two stimuli are shown sequentially, separated by a delay period. The duration of each stimulus is randomly sampled on each trial. The subject needs to judge which stimulus has a longer duration, and reports its decision during the decision period by choosing one of the two choice options.

Motor timing task (J. Wang et al. 2018)
Agents have to produce different time intervals using different effectors (actions).

Multi-sensory integration task

Two stimuli are shown in two input modalities. Each stimulus points to one of the possible responses with a certain strength (coherence). The correct choice is the response with the highest summed strength from both stimuli. The agent is therefore encouraged to integrate information from both modalities equally.

123-Go task (Egger et al. 2019)
Agents reproduce time intervals based on two samples.

Perceptual decision making task (Britten et al. 1992)
Two-alternative forced choice task in which the subject has to integrate two stimuli to decide which one is higher on average. A noisy stimulus is shown during the stimulus period. The strength (coherence) of the stimulus is randomly sampled every trial. Because the stimulus is noisy, the agent is encouraged to integrate the stimulus over time.

Delay response task (Inagaki et al. 2019)
Perceptual decision-making with delayed responses. Agents have to integrate two stimuli and report which one is larger on average after a delay.

Post-decision wager task (Kiani and Shadlen 2009)
Post-decision wagering task assessing confidence. The agent first performs a perceptual discrimination task (see for more details the PerceptualDecisionMaking task). On a random half of the trials, the agent is given the option to abort the sensory discrimination and to choose instead a sure-bet option that guarantees a small reward. Therefore, the agent is encouraged to choose the sure-bet option when it is uncertain about its perceptual decision.

Probabilistic reasoning task (T. Yang and Shadlen 2007)
The agent is shown a sequence of stimuli. Each stimulus is associated with a certain log-likelihood of the correct response being one choice versus the other. The final log-likelihood of the target response being, for example, option 1, is the sum of all log-likelihood associated with the presented stimuli. A delay period separates each stimulus, so the agent is encouraged to learn the log-likelihood association and integrate these values over time within a trial.

Pulse decision making task (Scott et al. 2015)
Pulse-based decision making task. Discrete stimuli are presented briefly as pulses.

Reaching task (Georgopoulos, Schwartz, and Kettner 1986)
Reaching to the stimulus. The agent is shown a stimulus during the fixation period. The stimulus encodes a one-dimensional variable such as a movement direction. At the end of the fixation period, the agent needs to respond by reaching towards the stimulus direction.

Reaching with self distraction task
In this task, the reaching state itself generates strong inputs that overshadows the actual target input. This task is inspired by behavior in electric fish where the electric sensing organ is distracted by discharges from its own electric organ for active sensing. Similar phenomena in bats.

Reaching task with a delay period task

A reaching direction is presented by the stimulus during the stimulus period. Followed by a delay period, the agent needs to respond to the direction of the stimulus during the decision period.

Ready-Set-Go task (Remington et al. 2018)
Agents have to measure and produce different time intervals. A stimulus is briefly shown during a ready period, then again during a set period. The ready and set periods are separated by a measure period, the duration of which is randomly sampled on each trial. The agent is required to produce a response after the set cue such that the interval between the response and the set cue is as close as possible to the duration of the measure period.

Spatial suppression motion task (Tadin et al. 2003)
This task is useful to study center-surround interaction in monkey MT and human psychophysical performance in motion perception. Tha task is derived from (Tadin et al. 2003). In this task, there is no fixation or decision stage. We only present a stimulus and a subject needs to perform a 4-AFC motion direction judgement. The ground-truth is the probabilities for choosing the four directions at a given time point. The probabilities depend on stimulus contrast and size, and the probabilities are derived from emprically measured human psychophysical performance. In this version, the input size is 4 (directions) x 8 (size) = 32 neurons. This setting aims to simulate four pools (8 neurons in each pool) of neurons that are selective for four directions.

Appendix II. Wrappers

Monitor
Saves relevant behavioral information: rewards, actions, observations, new trial, ground truth.

Noise
Add Gaussian noise to the observations.

Pass action
Modifies observation by adding the previous action.

Pass reward
Modifies observation by adding the previous reward.

Reaction time
Modifies a given environment by allowing the network to act at any time after the fixation period.

SideBias
Changes the probability of ground truth.

References

Abbott, Larry F., Dora E. Angelaki, Matteo Carandini, Anne K. Churchland, Yang Dan, Peter Dayan, Sophie Deneve, et al. 2017. "An International Laboratory for Systems and Computational Neuroscience." *Neuron* 96 (6): 1213–18.

Barak, Omri, Misha Tsodyks, and Ranulfo Romo. 2010. "Neuronal Population Coding of Parametric Working Memory." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 30 (28): 9424–30.

Barbosa, Joao, Heike Stein, Rebecca L. Martinez, Adrià Galan-Gadea, Sihai Li, Josep Dalmau, Kirsten C. S. Adam, Josep Valls-Solé, Christos Constantinidis, and Albert Compte. 2020. "Interplay between Persistent Activity and Activity-Silent Dynamics in the Prefrontal Cortex Underlies Serial Biases in Working Memory." *Nature Neuroscience* 23 (8): 1016–24.

Britten, K. H., M. N. Shadlen, W. T. Newsome, and J. A. Movshon. 1992. "The Analysis of Visual Motion: A Comparison of Neuronal and Psychophysical Performance." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 12 (12): 4745–65.

Brockman, Greg, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. "OpenAI Gym." *arXiv [cs.LG]*. arXiv. http://arxiv.org/abs/1606.01540.

Carnevale, Federico, Victor de Lafuente, Ranulfo Romo, Omri Barak, and Néstor Parga. 2015. "Dynamic Control of Response Criterion in Premotor Cortex during Perceptual Detection under Temporal Uncertainty." *Neuron* 86 (4): 1067–77.

Chaisangmongkon, Warasinee, Sruthi K. Swaminathan, David J. Freedman, and Xiao-Jing Wang. 2017. "Computing by Robust Transience: How the Fronto-Parietal Network Performs Sequential, Category-Based Decisions." *Neuron* 93 (6): 1504–17.e4.

Cueva, Christopher J., and Xue-Xin Wei. 2018. "Emergence of Grid-like Representations by Training Recurrent Neural Networks to Perform Spatial Localization." *arXiv [q-bio.NC]*. arXiv. http://arxiv.org/abs/1803.07770.

Daw, Nathaniel D., Samuel J. Gershman, Ben Seymour, Peter Dayan, and Raymond J. Dolan. 2011. "Model-Based Influences on Humans' Choices and Striatal Prediction Errors." *Neuron* 69 (6): 1204–15.

Egger, Seth W., Evan D. Remington, Chia-Jung Chang, and Mehrdad Jazayeri. 2019. "Internal Models of Sensorimotor Integration Regulate Cortical Dynamics." *Nature Neuroscience* 22 (11): 1871–82.

Farrell, Matthew, Stefano Recanatesi, Timothy Moore, Guillaume Lajoie, and Eric Shea-Brown. 2019. "Recurrent Neural Networks Learn Robust Representations by Dynamically Balancing Compression and Expansion." *bioRxiv*. https://doi.org/10.1101/564476.

Feulner, Barbara, and Claudia Clopath. 2021. "Neural Manifold under Plasticity in a Goal Driven Learning Behaviour." *PLoS Computational Biology* 17 (2): e1008621.

Freedman, David J., and John A. Assad. 2006. "Experience-Dependent Representation of Visual Categories in Parietal Cortex." *Nature* 443 (7107): 85–88.

Genovesio, Aldo, Satoshi Tsujimoto, and Steven P. Wise. 2009. "Feature- and Order-Based Timing Representations in the Frontal Cortex." *Neuron* 63 (2): 254–66.

Georgopoulos, A. P., A. B. Schwartz, and R. E. Kettner. 1986. "Neuronal Population Coding of Movement Direction." *Science* 233 (4771): 1416–19.

Hochreiter, Sepp, and Jürgen Schmidhuber. 1997. "Long Short-Term Memory." *Neural Computation* 9 (8): 1735–80.

Inagaki, Hidehiko K., Lorenzo Fontolan, Sandro Romani, and Karel Svoboda. 2019. "Discrete Attractor Dynamics Underlies Persistent Activity in the Frontal Cortex." *Nature* 566 (7743): 212–17.

Kiani, Roozbeh, and Michael N. Shadlen. 2009. "Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex." *Science* 324 (5928): 759–64.

Kobak, Dmitry, Wieland Brendel, Christos Constantinidis, Claudia E. Feierstein, Adam

Kepecs, Zachary F. Mainen, Xue-Lian Qi, Ranulfo Romo, Naoshige Uchida, and Christian K. Machens. 2016. "Demixed Principal Component Analysis of Neural Population Data." *eLife* 5 (April). https://doi.org/10.7554/eLife.10989.

Mante, Valerio, David Sussillo, Krishna V. Shenoy, and William T. Newsome. 2013. "Context-Dependent Computation by Recurrent Dynamics in Prefrontal Cortex." *Nature* 503 (7474): 78–84.

Masis, J. A., T. Chapman, J. Y. Rhee, D. D. Cox, and A. M. Saxe. 2020. "Rats Strategically Manage Learning during Perceptual Decision Making." *bioRxiv*. https://www.biorxiv.org/content/10.1101/2020.09.01.259911v1.abstract.

Mastrogiuseppe, Francesca, and Srdjan Ostojic. 2018. "Linking Connectivity, Dynamics, and Computations in Low-Rank Recurrent Neural Networks." *Neuron* 99 (3): 609–23.e29.

Miller, E. K., C. A. Erickson, and R. Desimone. 1996. "Neural Mechanisms of Visual Working Memory in Prefrontal Cortex of the Macaque." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 16 (16): 5154–67.

Mnih, Volodymyr, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. "Asynchronous Methods for Deep Reinforcement Learning." In *Proceedings of The 33rd International Conference on Machine Learning*, edited by Maria Florina Balcan and Kilian Q. Weinberger, 48:1928–37. Proceedings of Machine Learning Research. New York, New York, USA: PMLR.

Molano-Mazon, Manuel, Daniel Duque, Guangyu Robert Yang, and Jaime de la Rocha. 2021. "Pre-Training RNNs on Ecologically Relevant Tasks Explains Sub-Optimal Behavioral Reset." https://doi.org/10.1101/2021.05.15.444287.

Munoz, Douglas P., and Stefan Everling. 2004. "Look Away: The Anti-Saccade Task and the Voluntary Control of Eye Movement." *Nature Reviews. Neuroscience* 5 (3): 218–28.

Padoa-Schioppa, Camillo, and John A. Assad. 2006. "Neurons in the Orbitofrontal Cortex Encode Economic Value." *Nature* 441 (7090): 223–26.

Paszke, Adam, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, et al. 2019. "PyTorch: An Imperative Style, High-Performance Deep Learning Library." In *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d\textquotesingle Alché-Buc, E. Fox, and R. Garnett. Vol. 32. Curran Associates, Inc. https://proceedings.neurips.cc/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf.

Remington, Evan D., Devika Narain, Eghbal A. Hosseini, and Mehrdad Jazayeri. 2018. "Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics." *Neuron* 98 (5): 1005–19.e5.

Roitman, Jamie D., and Michael N. Shadlen. 2002. "Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 22 (21): 9475–89.

Rose, Nathan S., Joshua J. LaRocque, Adam C. Riggall, Olivia Gosseries, Michael J. Starrett, Emma E. Meyering, and Bradley R. Postle. 2016. "Reactivation of Latent Working Memories with Transcranial Magnetic Stimulation." *Science* 354 (6316): 1136–39.

Sarafyazd, Morteza, and Mehrdad Jazayeri. 2019. "Hierarchical Reasoning by Neural Circuits in the Frontal Cortex." *Science* 364 (6441). https://doi.org/10.1126/science.aav8911.

Saxena, S., A. Russo, J. Cunningham, and M. M. Churchland. 2021. "Motor Cortex Activity across Movement Speeds Is Predicted by Network-Level Strategies for Generating Muscle Activity." *bioRxiv*. https://www.biorxiv.org/content/10.1101/2021.02.01.429168v1.abstract.

Schaeffer, Rylan, Mikail Khona, Leenoy Meshulam, Ila Rani Fiete, and International Brain Laboratory. 2020 "Reverse-Engineering Recurrent Neural Network Solutions to a Hierarchical Inference Task for Mice." https://doi.org/10.1101/2020.06.09.142745.

Scott, Benjamin B., Christine M. Constantinople, Jeffrey C. Erlich, David W. Tank, and Carlos D. Brody. 2015. "Sources of Noise during Accumulation of Evidence in Unrestrained and Voluntarily Head-Restrained Rats." *eLife* 4 (December): e11308.

Steinmetz, Nicholas A., Peter Zatka-Haas, Matteo Carandini, and Kenneth D. Harris. 2019. "Distributed Coding of Choice, Action and Engagement across the Mouse Brain." *Nature* 576 (7786): 266–73.

Sussillo, David, Mark M. Churchland, Matthew T. Kaufman, and Krishna V. Shenoy. 2015. "A Neural Network That Finds a Naturalistic Solution for the Production of Muscle Activity." *Nature Neuroscience* 18 (7): 1025–33.

Sutton, Richard S., and Andrew G. Barto. 2018. *Reinforcement Learning, Second Edition: An Introduction*. MIT Press.

Tadin, Duje, Joseph S. Lappin, Lee A. Gilroy, and Randolph Blake. 2003. "Perceptual Consequences of Centre–surround Antagonism in Visual Motion Processing." *Nature* 424 (6946): 312–15.

Wang, Jane X., Zeb Kurth-Nelson, Dharshan Kumaran, Dhruva Tirumala, Hubert Soyer, Joel Z. Leibo, Demis Hassabis, and Matthew Botvinick. 2018. "Prefrontal Cortex as a Meta-Reinforcement Learning System." *Nature Neuroscience* 21 (6): 860–68.

Wang, Jing, Devika Narain, Eghbal A. Hosseini, and Mehrdad Jazayeri. 2018. "Flexible Timing by Temporal Scaling of Cortical Responses." *Nature Neuroscience* 21 (1): 102–10.

Werbos, P. J. 1990. "Backpropagation through Time: What It Does and How to Do It." *Proceedings of the IEEE. Institute of Electrical and Electronics Engineers* 78 (10): 1550–60.

Yang, Guangyu Robert, and Manuel Molano-Mazon. 2021. "Next-Generation of Recurrent Neural Network Models for Cognition." *PsyArXiv*. https://doi.org/10.31234/osf.io/w34n2.

Yang, Tianming, and Michael N. Shadlen. 2007. "Probabilistic Reasoning by Neurons." *Nature* 447 (7148): 1075–80.

Zhang, Xiaoxing, Wenjun Yan, Wenliang Wang, Hongmei Fan, Ruiqing Hou, Yulei Chen, Zhaoqin Chen, et al. 2019. "Active Information Maintenance in Working Memory by a Sensory Cortex." *eLife* 8 (June). https://doi.org/10.7554/eLife.43191.