

Evolution of Biologically Plausible Neural Networks Performing a Visually Guided Reaching Task

Derrick E. Asher

Jeffrey L. Krichmar

Nicolas Oros

Department of Cognitive Sciences
2201 Social & Behavioral Sciences Gateway Building
University of California, Irvine CA 92697-5100
{dasher, jkrichma, noros}@uci.edu

ABSTRACT

An evolutionary strategy (ES) algorithm was utilized to evolve a simulated neural network based on the known anatomy of the posterior parietal cortex (PPC), to perform a visually guided reaching task. In this task, a target remained visible for the duration of a trial, and an agent's goal was to move its hand to the target as rapidly as possible and remain for the duration of that trial. The ES was used to tune the strength of 15609 connections between neural areas and 4 parameters governing the neural dynamics. The model had sensory latencies replicating those found in recording studies with monkeys. The ES ran 100 times and generated very diverse networks that could all perform the task well. The evolved networks 1) showed velocity profiles consistent with biological movements, and 2) found solutions that reflect short-range excitation and long-range, contralateral inhibition similar to neurobiological networks. These results provide theoretical evidence for the important parameters and projections governing sensorimotor transformations in neural systems.

General Terms

Algorithms, Experimentation, Theory

Keywords

Neural Networks, Visually Guided Reaching, Evolutionary Strategy, Sensorimotor Transformation

1. INTRODUCTION

Primates have the impressive ability to perform fluid movements guided by visual and proprioceptive information. Complex neural architectures in the primate brain that are responsible for these sensory guided movements have emerged through the process of evolution and natural selection. In particular, experimental evidence has highlighted specific roles for the parietal, pre-motor and motor cortices. The parietal cortex has been viewed as an association area that combines different types of sensory information to form a single representation of space [1]. Evidence has implicated the posterior parietal cortex (PPC) as being necessary for transforming multiple sources of sensory information such as visual and proprioceptive information into actions [2, 3], and important for aspects of movement [4, 5]. Sensory integration for the control of reaching actions is primarily performed by a network of brain regions connecting the PPC,

dorsal pre-motor (PMd) and, primary motor (M1) cortices [3, 6]. Interestingly, experiments have shown that the sensory inputs associated with the neural activity of reaching actions arrive in parietal brain regions with different latencies. It was found that proprioceptive inputs indicating the position of hand relative to the center of gaze arrives after a delay of approximately 30ms, and visual inputs from a stimulus arrive after a delay of approximately 90ms [7].

Computational models have contributed significantly towards our understanding of the neurophysiological activation associated with action execution and sensory-motor transformations [8-12]. Visually guided models of reaching actions have been used to show how sensory input combines to give rise to joint angle estimations [13], and the neural representations in distinct regions of the brain for competing reaching decisions [8]. However, the aforementioned models did not take into consideration the different sensory delays that must be accounted for in order to assemble a biologically realistic neural network model performing sensorimotor transformations. Furthermore, while they included feedforward projections from sensory to motor areas, they did not take into consideration the huge number and diversity of lateral connections within the PPC and feedback connections to the PPC.

Therefore, the main goal of the present work was to investigate if a neuroanatomical architecture with complex lateral and feedback connections could cope with sensory delays while performing a visually guided reaching task. The second goal of this work was to investigate the diversity of networks with the same general architecture but different parameter values that could perform the task well. Specifically, we used an evolutionary strategy (ES) algorithm to tune the large number of parameters (~15000) of our model. We ran the ES 100 times with different random seeds and analyzed the differences between the best solutions generated.

In the present work, we show that a biologically plausible neural architecture could be artificially evolved to perform accurate reaching movements to visual targets, and overcome sensory delays. The results suggest a neural architecture based on the parietal cortex for reaching, and provide theoretical evidence for the diversity of connection sets that can give rise to the behavior observed in humans and non-human primates. Such an architecture inspired by the primate brain highlights the great potential for neural based sensorimotor systems that can generate fluid sensory-guided movements.

2. METHODS

2.1 Visually Guided Reaching Task

Simulated agents were evolved through artificial selection to perform a visually guided task designed to test smooth and accurate movements in the presence of visual stimulus input (target). This task was designed to be similar to those carried out by humans and non-human primates [14]. An agent consisted of a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

GECCO '14, July 12 - 16 2014, Vancouver, BC, Canada

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2662-9/14/07...\$15.00.

<http://dx.doi.org/10.1145/2576768.2598368>

hand that could move on a 2D space, an eye that perceived the hand and the target in this space, and a neural network that processed sensory information into motor commands in order to move the hand towards the target. The 2D space was measured in degrees of visual angle from an agent’s central fixation ($0^\circ, 0^\circ$) ranging from $[-50^\circ, 50^\circ]$ in both the x and y directions. An agent’s vision was fixed, implying that it could not move its eyes or head.

During a trial, the goal of an agent was to move its hand towards a target as fast as possible, and then hold it on the target for the remainder of the trial. An agent always started a trial with its hand located in the center of the 2D space ($0^\circ, 0^\circ$). The target was visible to the agent for the entire duration of a trial (50 timesteps simulating 500ms of real time).

Agents were evolved to reach eight targets situated at different locations on an invisible square. For this reason, each agent performed eight trials, one trial for each of the targets located 25° (two vertical and two horizontal targets) or 35° (four diagonal targets) from the center of the visual space.

2.2 Neural Network Architecture

A neural network based on the known anatomy of the primate brain was designed to perform a visually guided reaching task. The model consisted of 367 neurons distributed into four distinct regions (Figure 1): Vision, Proprioception, Posterior Parietal Cortex (PPC), and Premotor/Primary Motor Cortex (PMd/M1) with a fixed number of neurons per region. Excitatory topographic projections were made between the Vision area and the PPC, and inhibitory topographic projections between the Proprioception area and PPC (Figure 1 and Section 2.2.1). The PPC had lateral (recurrent) connections where neurons were fully connected to each other. The PPC was also fully connected to PMd/M1 neurons encoding movement of the hand (Up, Down, Left, and Right). Finally, these neurons projected back to the neurons in the PPC.

In our experiments, only the weights from the PPC to PMd/M1 (feedforward - 484 weights), from PMd/M1 to PPC (feedback - 484 weights), and of the lateral PPC connections were evolved (14641 weights). Therefore, the model had a total of 15609 evolvable connections, initialized randomly to values between $[-1, 1]$ drawn from a uniform distribution.

For purposes of analyses, the parameters were separated into two distinct categories, neural firing and weights. The evolved parameters governing simulated neural firing rates (Section 2.2.2) were compared directly to determine significant relationships within and between neural layers. The evolved connections were first separated into three distinct weight sets (feedforward, feedback, and lateral) and compared accordingly. The feedforward and feedback weights were further separated into ipsilateral and contralateral. The ipsilateral weights consisted of connections between the PPC and PMd/M1 layers corresponding to the same side of the agent (e.g., left half of PPC neurons connected to the left PMd/M1 neuron), and the contralateral weights for the opposite side of the agent. Lateral connections were further separated into short, medium, and long-range weights determined by the distance between PPC neurons.

2.2.1 Sensory Inputs

An agent received visual information from both the hand and the target, along with proprioceptive input from the hand. Both Vision and Proprioception regions of the network consisted of 11×11 2D arrays of sensory neurons, spanning the space $[-50^\circ, 50^\circ]$ with a resolution of 10° . The sensory information from the locations of the hand and the target was mapped into the activity of the

corresponding neurons in the Proprioception and Vision areas (activity set to one). The activity of the other sensory neurons was set to zero. In the case of vision, if both the hand and target overlapped, the activity of the visual input neuron for that particular location was set to two.

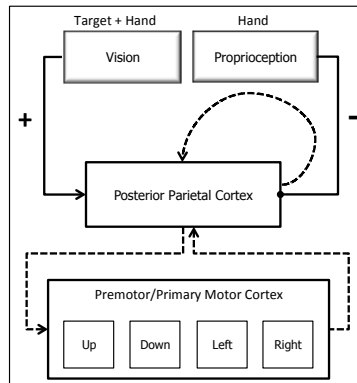


Figure 1. Neural architecture consisting of four areas: vision, proprioception, posterior parietal cortex (PPC) and premotor/primary motor cortex (PMd/M1). Dashed lines indicate evolvable connections.

The Vision to PPC connections were solely excitatory, whereas the Proprioception to PPC connections were solely inhibitory, allowing these sensory inputs to combine in a subtractive manner, similar to the empirically derived vector subtraction that seems to be computed in the PPC [15]. The weight values of the projections from the Vision and Proprioception areas to the PPC neurons were assigned using 2D cosines as defined by Equation 1:

$$w_{ij} = f * \cos^g \left(\frac{d}{20} \right) \quad (1)$$

where w_{ij} is the weight value connecting neuron j to neuron i , f is a scale factor set to 2 for vision and -4 for proprioception. The value of 2 was selected as the amplitude for the visual input signal. The value -4 was selected for the proprioceptive amplitude to approximately cancel the visual signal when the hand covered the target. The distance d was measured between a sensory and a PPC neuron. The peak of the 2D cosine resided at the location where a sensory neuron’s coordinates were the same as a PPC neuron’s coordinates ($d = 0$). The gain g was set to 200 to ensure that the projections from Proprioception and Vision to PPC displayed broad cosine tuning curves across multiple PPC neurons.

Based on empirical evidence [16], we implemented sensory latencies, which provided our model with delayed information on the location of the hand. It was observed that proprioceptive inputs have a latency of at least 30ms and visual inputs have a latency of about 90ms. To implement these latencies, the PPC neurons were provided with proprioceptive spatial information about the hand’s location that lagged by three timesteps and visual information that lagged by nine timesteps. This implies that the model did not receive any visual input from the target or hand locations until the 9th timestep of every trial, rendering the model blind to the task for the first $1/5^{\text{th}}$ of each trial. Proprioceptive information was received upon the 3rd timestep of each trial, though it provided no information about the target’s location.

2.2.2 PPC and PMd/M1 Neural Firing

Neurons within the PPC and PMd/M1 regions of the model generated firing rates calculated with a sigmoid function that

bounded activity between [0, 1]. The sigmoid function is described in Equation 2:

$$s_i(t) = \frac{1}{1 + \exp(\Delta - I_i(t)) * k} \quad (2)$$

where $s_i(t)$ is the firing rate for neuron i at timestep t , Δ is the bias for the sigmoid function that shifts the range of sensitivity, $I_i(t)$ is the synaptic input into neuron i from all connected neurons at timestep t , and k is the gain on the sigmoid function, which determines the sensitivity of the neurons. Neurons from the same area shared the same bias and gain values, but different areas could have different values. The synaptic input $I_i(t)$ was given by:

$$I_i(t) = \sum_j w_{ij} s_j(t-1) \quad (3)$$

where $I_i(t)$ is the input to a PPC or PMd/M1 neuron i at timestep t , w_{ij} is the weight connecting from neuron j to neuron i , and $s_j(t-1)$ is the firing rate for neuron j at the previous timestep ($t-1$).

These bias and gain parameters were evolved for the PPC and PMd/M1 areas using the evolutionary strategy described below. Therefore, four neural firing parameters were evolved (two for each area) with the values of Δ and k being continuous and bounded within the ranges [-5, 5] and [0.1, 10] respectively.

2.3 Hand Movement

The reaching movements of the hand in the 2D space were calculated by a velocity vector defined by a rate of change in position (visual degrees/second) and a direction (angle). First we multiplied the activity of the PMd/M1 neurons with the sine and cosine of their respective angles in polar coordinates (Up: $\pi/2$, Left: π , Down: $3\pi/2$, and Right: 2π). The vector sum of these values was calculated and used to compute the horizontal and vertical components of the velocity vector. This method is commonly referred to as population coding [17]. The population coding method is described in Equation 4:

$$D_i = v * \sum_j \sin(\theta_j - \varphi_i) * s_j(t) \quad \begin{cases} i = 1, 2 \\ j = 1, \dots, 4 \end{cases}$$

$$\theta_j = (j - 1) * \frac{\pi}{2}$$

$$\varphi_i = (i - 1) * \frac{\pi}{2} \quad (4)$$

where D is the velocity vector defined by vertical ($i=1$) and horizontal ($i=2$) components, j is the index representing the four PMd/M1 neurons, θ_j is the angle representing right, up, left, and down respectively, φ_i shifts the angle θ by $\pi/2$ to take the cosine for the horizontal component of the velocity vector D , $s_j(t)$ is the firing rate for a PMd/M1 neuron j at timestep t , and v is a scalar set to 2. This value was chosen to set the maximum velocity of the hand to 2 degrees per timestep (200 degrees/second) if moving to a horizontal or vertical target (e.g., $D_1 = 2.0$ and $D_2 = 0$), and 2.83 degrees per timestep (283 degrees/second) if moving towards a diagonal target (e.g., $D_1 = 2.0$ and $D_2 = 2.0$). This rate of movement was chosen to match biological movement recorded in human and non-human primate reaching studies [18, 19].

2.4 Evolutionary Strategy

An evolutionary strategy algorithm, implemented in the Evolving Objects (<http://eodev.sourceforge.net>) library [20], was used to tune the neural networks over 25000 generations. At that point, the agent with the lowest fitness value was selected as the best

agent. The free parameters and their ranges of values were: 1) 15609 connection weights [-1.0, 1.0] (Figure 1; dashed arrows), two biases [-5.0, 5.0] (Δ in Equation 2) and two gains [0.1, 10] (k in Equation 2) for PMd/M1 and PPC areas (15613 total free parameters). All these parameters were continuous floats causing the search space to be extremely large. The ES used in our experiments was the evolution strategy algorithm named ‘‘mu comma lambda’’ (μ, λ) that includes elitism, crossover with roulette wheel selection, and mutation [21] with a population size of 20 agents. Each agent had a 40% chance of undergoing mutations. Mutations were applied randomly (probability of 0.5) to all free parameters by the addition of a small random number drawn from a Gaussian distribution $G(0, 0.3)$ for the weights, $G(0, 3)$ for the biases, and $G(0, 1.5)$ for the gains. The same ES was run 100 times with different random seeds to produce 100 independently evolved best agents. For the remainder of the manuscript the 100 independently evolved best agents will be referred to as the population.

2.4.1 Fitness Function

The fitness function was defined to reflect both an agent’s speed and accuracy when reaching to the eight targets, as well as the ability to stay on the targets. Lower fitness values indicated better performance. Fitness was calculated by summing the Euclidian distance from the agent’s hand to the target at each timestep, across all trials. The fitness function is given by Equation 5:

$$F = \sum_n \sum_t E_n(t) \quad \begin{cases} n = 1, \dots, 8 \\ t = 1, \dots, 50 \end{cases}$$

$$E_n(t) = \sqrt{(H_x(t) - T_{xn})^2 + (H_y(t) - T_{yn})^2} \quad (5)$$

where F is an agent’s fitness value, n represents the trial, t is the timestep in a trial, $E_n(t)$ is the Euclidian distance between an agent’s hand and the target for timestep t during trial n , $(H_x(t), H_y(t))$ are the x and y components of hand position in 2D Cartesian space, and (T_{xn}, T_{yn}) is the 2D Cartesian coordinate of the target position for trial n .

The best (minimum) possible fitness value was 3563.4, which was the total amount of degrees of visual angle across all timesteps and trials that an agent would take to move to the eight targets at maximum velocity. Therefore, we subtracted off 3563.4 from all agents’ fitness values to provide a normalized fitness per agent that was optimized at 0.

3. RESULTS

3.1 Fitness Evolution

Fitness values from the 100 agents generated by 100 ES runs show that fitness improved exponentially for the first 5000 generations and linearly from approximately generation 5000 to 25000 (Figure 2). It is clear from Figure 2 that the best agent evolved in coincidence with the population until generation 2000 (Figure 2; dashed line). From generations 2000 to 25000, the same ES consistently evolved a better fitness than the other 99 ES runs, resulting in a final fitness of 519.7 at generation 25000. The solid black line in Figure 2 is the mean fitness for the population of evolved agents with the width of the line showing the standard deviation around the mean for every generation, with a final average fitness of 1563.5 ± 363.9 at generation 25000.

3.2 Agent Behavior

To demonstrate that the agents evolved correct behaviors, we recorded the reach paths, velocity profiles, and target error (TE)

for each of the 100 agents. The reach paths show the reaching trajectories across the trials, indicating precisely how an agent moved towards targets. The velocity profiles are complementary to the reach paths, in that they are derived from the reach path data, yet they provide a better understanding of why one agent’s fitness was better than another. The TE shows how close an agent was to the target upon the last timestep of every trial, which was independent of an agent’s reach path or velocity profile.




Figure 2. Fitness evolution over generations. The y-axis shows the fitness values. The x-axis shows generations. The thick solid line is the average fitness for the population. Dashed line shows the ES that produced the best agent.

3.2.1 Reach Paths

Reach paths for an agent’s hand position at every timestep are shown in Figure 3 (a-c). Each data point represents either an agent’s absolute hand position in space per timestep (Figure 3a-b), or the average hand position across the population per timestep (Figure 3c), with each line representing a trial. Distances between data points indicate the rate of movement.

Agents’ trajectories in the visually guided reaching task reflect excellent performance corresponding to fitness values and distance from the target at the end of each trial (TE). We plotted the path of two individual agents: the best agent (agent #1) and the worst agent (agent #100). The mean reach path for the population of agents was calculated for every timestep across all agents and trials (Figure 3c). A comparison between the two individual agents illustrates one aspect of the diversity that emerged as a result of running 100 ES. The reach paths for agent #1 are clearly shorter (straighter) than that of agent #100. However, it can be seen that even agent #100 was able to correctly perform the reaching task. This point is emphasized with the average reach path from the population (Figure 3c).

The agents’ reach paths provide a good understanding of the different reaching movements that emerged, but do not emphasize the rates at which the agents moved. We therefore investigated the velocity profiles of the agents.

3.2.2 Velocity Profiles

Figure 3 (d-f) shows the velocity profile for agent #1, agent #100, and the population. These were plotted using the average rate of movement per timestep (instantaneous rate) across the eight trials. It is important to note that the velocity profiles match up with the reach paths (Figure 3a-3c), implying that the velocity profiles (Figure 3; bottom row) reflect the average rate at which the hand moved through space in the reach path plots (Figure 3; top row).

The velocity profiles shown in Figure 3 (d-f) demonstrate that the agents’ acceleration reached maximum velocity followed by

deceleration until movement ceased. This data is similar to that observed from empirical studies with control subjects [18, 19].

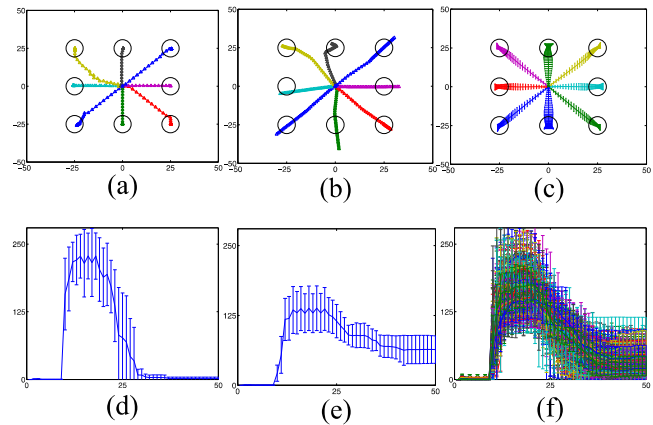


Figure 3. Reach paths and velocity profiles. (a-b) Reach paths for agent #1 and agent #100. (c) Average reach path for the 100 agents. (d-e) Average velocity profile for agent #1 and #100. (f) Overlapped average velocity profile for all agents. Error bars are standard error of the mean (SEM).

There are multiple differences between the two individual agents’ data shown. Agent #1 (Figure 3d) moved its hand at a fast average rate then ceased movement at about 30 timesteps. Agent #100 moved about half as fast as agent #1 and had difficulty stopping. Figure 3f suggests that the population of agents were able to efficiently deal with sensory delays by the emergence of a single optimized movement towards the intended goal. A single-hump in the velocity profile suggests that the evolved architecture (Figure 1) allowed for sensory integration (proprioceptive and visual) of the hand positions through the trials, resulting in precise reaches to the visual targets (see Figure 3). Taken together, these results show that the population produced realistic behaviors.

3.2.3 Target Error (TE)

TE is a complementary measure of behavioral performance, which reveals the accuracy of the agents, independent of movement rate or path. TE corresponds to the distance between an agent and the target at the last timestep of each trial. Figure 4a shows that #100 was more accurate than #1 on half of the trials.

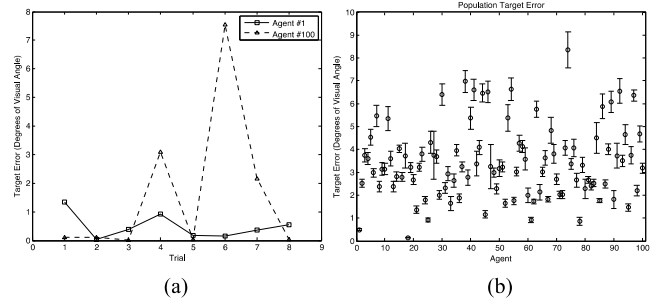


Figure 4. Target error (TE). Distance between target and hand at the last timestep per trial. Error bars show standard error of the mean (SEM). (a) TE across trials for agents #1 and #100. (b) Average TE per agent.

Interestingly, it is not clear from the TE metric that agent #1 has performed better than agent #100 across trials. This is due to the fact that TE does not take into account the rate of movement (velocity profile) or the distance (reach path) taken to reach the target. However, the TE metric does show how diverse behaviors were between the two agents across trials (Figure 4a). Moreover,

average TE across the population reflects differences in behavior that independently evolved (Figure 4b).

Interestingly, we observed that the average TE for each agent was below the implemented 10° sensory (visual and proprioceptive) resolution. The TE metric emphasizes the fact that the evolved agents showed great accuracy despite the limited visual resolution. Further experiments will be conducted in order to investigate this emergent property.

3.3 Parameter Analysis

3.3.1 Population Diversity

We investigated the diversity of evolved parameters in the population of solutions to confirm that they were independent and non-degenerate. We used the moment of inertia (MoI) for quantifying population diversity of different parameters sets [22]. Values closer to 0 indicate less diversity, whereas greater values represent greater population diversity. Through the evolutionary process, diversity often decreases as the algorithm settles on a solution. Therefore, it is common practice to show the MoI in early vs. late stages of evolution to provide an understanding of how much diversity vanished as evolution progressed.

The amount of parameters along with standard deviation around an average parameter, the range of possible parameter values and the moment of inertia (MoI) calculation provide a decent approximation of population diversity (Table 1). The values in the rows of Table 1 are complementary, indicating that they are similar ways of assessing the diversity of a population. The amount of continuous parameters (Table 1; # Parameters) and their respective ranges (Table 1; Range) indicate the size of the search space that the solution(s) reside in. Given that the solution(s) are good, the average parameter values and their respective deviations indicate the relevant solution values and their ranges, provided that the population is diverse. Therefore, alternative metrics like MoI comparisons provide reinforcing evidence about population diversity.

Table 1. Population diversity measurements. Number of parameters, average parameter value (stdev), range of parameter values, and moment of inertia (MoI).

Parameter	# Parameters	Range	Mean (\pm s.d.)	MoI _[Gen 1]	MoI _[Gen 25000]
PPC Bias	1	[-5.0, 5.0]	1.831 \pm 0.582	932.536	33.538
PPC Gain	1	[0.1, 10.0]	5.670 \pm 2.483	912.125	610.539
PMd/M1 Bias	1	[-5.0, 5.0]	0.909 \pm 0.852	965.219	72.917
PMd/M1 Gain	1	[0.1, 10.0]	1.906 \pm 1.697	838.444	284.948
Weights: feedforward	484	[-1.0, 1.0]	0.006 \pm 0.582	33.897	32.082
Weights: feedback	484	[-1.0, 1.0]	0.006 \pm 0.577	33.673	32.972
Weights: lateral	14641	[-1.0, 1.0]	0.002 \pm 0.577	33.978	32.997

In order to assess how much each evolved parameter contributed to the diversity of the population, the MoI was calculated for each parameter listed in Table 1 individually. It can be seen that the weights contributed the least to population diversity and showed minimal change after 25000 generations of evolution (Table 1; MoI_[Gen 1] to MoI_[Gen 25000]). Weights: feedforward, feedback, and lateral). In contrast, the evolved gain and bias values contributed the most to diversity of the evolved population. Of these four parameters, PPC bias contributed the least to final population diversity, but had the greatest change in MoI (898.998), indicating a narrow range of solutions found by the ES, potentially indicating the importance of this parameter in good solutions. The PPC gain contributed the most to final population diversity, but had the smallest change in MoI (301.586), implying the flexibility of this parameter in the solutions. The PMd/M1 bias MoI was roughly double the PPC bias, but small relative to the MoI of the PPC gain, indicating that the narrow range of this evolved

parameter could also be important for good solutions. The PMd/M1 gain, like the PPC gain contributed a large portion of the final population diversity, leaving the relative importance of this parameter in between the PMd/M1 bias and the PPC gain.

Together, the ES found a diverse set of good solutions indicated by the parameter metrics shown in Table 1 paired with the behavioral data shown in Figures 3-4. The population diversity metrics across the parameters provide insight into the potential importance of select parameters and their relative contributions to population diversity across the set of evolved solutions.

3.3.2 Neuron Parameters

To evaluate the relationship between parameters, pairwise linear sample correlation coefficients were calculated (Table 2). The significant comparisons (Table 2; bold red values) reflect the Pearson's correlation coefficient r , which indicates the strength and slope (positive or negative) of the relationship. The four evolved neuron parameters and fitness (Table 2) were considered in the pairwise linear comparisons to determine how the parameters influenced each other, and which parameters influenced fitness. The considered parameters were: PPC bias, PPC gain, PMd/M1 bias, PMd/M1 gain, and fitness. The majority of the pairwise comparisons did not show a linear relationship. Three linear comparisons showed a significant pairwise relationship: 1) PPC gain to PPC bias, 2) PMd/M1 bias to PPC bias, and 3) PMd/M1 gain to fitness. Alpha was 0.00036 Bonferroni corrected for multiple comparisons at the 99 percent confidence level. Indicating that these three parameter comparisons were highly significant ($p < 0.0001$), and reflected a strong negative linear relationship for the 100 evolved solutions.

Table 2. Pearson's sample correlation coefficients between pairs of evolved parameters and fitness. Sign indicates slope of the relationship and bold, red values indicate significance.

Parameter	PPC Bias	PPC Gain	PMd Bias	PMd Gain
PPC Bias	--	--	--	--
PPC Gain	-0.713	--	--	--
PMd Bias	-0.733	0.303	--	--
PMd Gain	0.129	-0.031	-0.1457	--
Fitness	0.260	-0.268	-0.068	-0.524

In order to further understand the relationship and differences between the parameters of the 100 solutions, a hierarchical cluster analysis was performed on the same parameters as shown in Table 2. The hierarchical clustering was calculated by taking the largest distance between data points, also called *furthest neighbor* or *complete-linkage clustering* (Figure 5a). Seven clusters were identified that represented a conservative hierarchical separation in agents' parameters (including fitness). The cluster threshold was selected to demonstrate conservative separation between parameter clusters. The cluster separation in the dendrogram (Figure 5a) was mapped directly to the pairwise parameter cluster comparisons shown in Figure 5(b-d). The colored polygons in the cluster plots simply connect the perimeter data points to outline each cluster for easier viewing. Parameter values for agents #1 and #100 (Figure 5) tie behavior to parameters.

The clusters reveal marked separations for the PPC gain vs. PPC bias (Figure 5b). Interestingly, the cluster of which the best solution (Figure 5b; #1) shares commonalities with six of the seven clusters, implicating that cluster as a range and domain for optimal PPC gain and PPC bias values. It is interesting to note that relative to the evolved population, solution #1 (best) evolved

a high PPC gain and low PPC bias compared to solution #100 (worst), which evolved the opposite, a high PPC bias and low PPC gain. However, although there exists a significant linear relationship between these evolved parameters (Table 2), a conclusion cannot be inferred about behavioral performance.

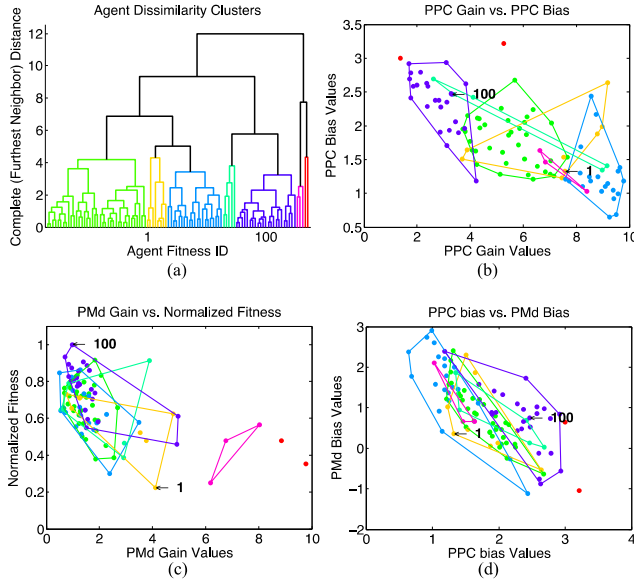


Figure 5. (a) Dendrogram of 7 hierarchical clusters based on evolved parameters and fitness for population. Cluster plots (b-d) show linear relationships between parameters. Enclosed points indicate clusters shown in dendrogram. (b) PPC gain vs. PPC bias. (c) Fitness vs. PMd/M1 gain. (d) PPC bias vs. PMd/M1 bias. Agents #1 and #100 are labeled in these plots.

In contrast to the diversity shown in the PPC parameters governing firing rates in the PPC layer, PMd/M1 gain vs. fitness (Figure 5c) revealed a large majority of solutions evolved PMd/M1 gain values between 0 and 2, which resulted in higher fitness values. Although similar to the yellow cluster overlap in Figure 5b, there is heavy overlap between five of the seven clusters in Figure 5c. However, the domain [0, 2] of PMd/M1 gain values did not lead to better fitness (Figure 5c). In fact, the data show that the top five agents (data points closest to 0 on the y-axis) evolved PMd/M1 gain values greater than 2, indicating that a PMd/M1 gain value greater than 2 was important for better fitness. Therefore, unlike the comparison between PPC gain and bias values, a PMd/M1 gain value greater than 2 led to better fitness, resulting in a better solution and behavior.

Interestingly, a significant linear relationship between the bias values in the two different layers (PPC and PMd/M1) evolved for the population (Figure 5d). The data show that the majority of solutions evolved PMd/M1 bias values between [0, 2] and PPC bias values between [1, 3]. This result suggests that these two parameters may have been critical for good solutions because of the narrow range and the significant linear relationship (Table 2). In addition, it is important to note that none of the evolved PPC bias values were below zero (all positive), whereas, several agents evolved negative PMd/M1 biases (Figure 5d). Furthermore, the overlapped clusters indicate that an optimal PMd/M1 and PPC bias value pair lies within a narrow range in 2-D Cartesian space.

Neuron parameters comparisons revealed three important aspects of the evolved parameters. 1) Population diversity remained after evolution (Table 1), indicating a diverse set of solutions found by the ES. 2) Significant linear relationships between evolved

parameters emerged (Table 2). 3) Parameter clusters pertained to different agent behavior and fitness (Figure 5).

3.3.3 Weight Parameters

The average population weights were analyzed to determine their contributions towards the observed behavior. In particular, we investigated the connection length (distance between nodes), sign (excitatory or inhibitory), and strength. The weight parameters consist of the three sets of connections, 1) projections from the 121 (11x11) PPC neurons to the 4 PMd/M1 neurons, 2) projections from the PMd/M1 neurons to the PPC neurons, and 3) recurrent projections between the PPC neurons. The first two sets of projections were analyzed with a comparison of the average connection strength of the 100 agents between half the PPC neurons and the corresponding PMd/M1 neuron (ipsilateral: e.g., upper (left) half of PPC neurons connecting to the *Up (Left)* PMd/M1 neuron), vs. the other half of the PPC neurons and the same PMd/M1 neuron (contralateral: e.g., lower (right) half of PPC neurons connecting to the




Figure 6. Analysis of evolved population weights. (a) Average ipsi vs. contra lateral feedforward weights (b) Average ipsi vs. contra lateral feedback weights. (c) Average short, medium and long lateral weights.

In contrast, recurrent PPC connections were analyzed by comparing connection length and average connection strength across the population of agents. Connection length was calculated with the Euclidian distance between each PPC neuron. The PPC connections were separated into short (distance < 5), medium (5 < distance < 8), and long (distance > 8) range projections. The average relative strength per length provided a metric for how the different length connections contributed to the evolved behavior.

Due to the topographic organization of the sensory input layers and their projections to the PPC, the population of agents evolved ipsilateral excitatory and contralateral inhibitory feedforward PPC to PMd/M1 projections (Figure 6a). Two sample Kolmogorov-Smirnov tests of the population average PPC to PMd/M1 ipsilateral excitatory and contralateral inhibitory projections revealed strong significance ($p \ll 0.0001$). This result indicates that ipsilateral PPC projections corresponding to the intended direction of movement (the direction of the visual target) excited the PMd/M1 neuron to move in that direction. Furthermore, if the visual stimulus appeared in the contralateral region of PPC neurons, the projections inhibited

that same PMd/M1 neuron, so movement towards the incorrect direction was inhibited.

In contrast to feedforward (PPC→PMd/M1) weights, feedback (PPC←PMd/M1) weights did not show significant differences when comparing ipsilateral to contralateral (Figure 6b). However, it is interesting to note that feedback ipsilateral and contralateral average projections were roughly the same (mean of zero), indicating that the feedback weights may not have contributed towards the observed behavior across the population of agents. Whereas, the ipsilateral vs. contralateral analysis revealed high significance in the feedforward projections and no significance in the feedback projections, the recurrent PPC length vs. average strength analysis resulted in a significant difference between short and long-range PPC projections (Figure 6c).

A two-sample Kolmogorov-Smirnov test was utilized to determine differences in weight values between the PPC short vs. medium, medium vs. long, and short vs. long-range weights. Neither the short vs. medium nor medium vs. long-range weight comparisons were significant, but short vs. long was significant ($p = 0.0134$; Bonferroni corrected for multiple comparisons) at the 0.05 confidence level. Furthermore, the significant difference between long and short-range weights across the recurrent PPC neurons indicates that on average, the population evolved short-range excitation and long-range inhibition. This result coincides with the significant difference found between feedforward ipsi and contralateral weights. These results help explain how the evolved weights contributed towards good task performance.

4. DISCUSSION

In this study, simulated agents were evolved through artificial selection to perform a visually guided task. We ran an evolutionary strategy algorithm 100 times to evolve a diverse set of biologically realistic neural networks. We found that 1) all 100 evolved agents successfully performed the task even though the parameter space was vast (15613 parameters), 2) the evolved parameters gave rise to velocity profiles consistent with visually guided reaching studies in control subjects [18, 19], 3) the population had high precision ($<5^\circ$) despite the implemented sensory resolution (10°), 4) linear relationships between specific neuron parameters emerged, 5) the evolved agents coped with the sensory inputs corresponding to different delays, 6) ipsilateral projections from PPC to PMd/M1 tended to be excitatory and contralateral tended to be inhibitory, and 7) recurrent PPC projections had short-range excitation and long-range inhibition.

In the field of artificial intelligence, evolutionary computation (EC) is often used to design or tune neural networks that control agents [23, 24]. EC has been used in computational neuroscience to fit the activity of neuron models to electrophysiological recordings [25, 26]. Most relevant to our work, neural networks were evolved to perform a reaching task with intentions of investigating the interaction between evolution and individual learning [27]. In this simulation study however, the sensory inputs and neural architectures were abstract and not intended to investigate the mammalian brain or cortical interactions. More recently, neural networks were evolved in robotic platforms to investigate the relationship between selective attention and action selection by running experiments in a foraging task [28]. Similar to our work, they investigated the relationship between evolved architecture and the resulting behavior, but their neural architectures were not intended to model sensory integration or sensorimotor transformations, nor did their models provide any predictions pertaining to the underlying biology. Furthermore, these implementations did not constrain their neural network connectivity

based on known neuroanatomy, consider empirically derived sensory delays, or attempt to use direct encoding for the evolution of a large parameter set (15613 parameters) that better reflects the intrinsic complexity of a brain network. Biologically plausible parieto-frontal neural networks performing sensorimotor transformations have been utilized to investigate the neural correlates between competing reaching decisions [8]. In these studies, the neural networks were hardwired, indicating that their connections were preset and unable to change. Additionally, the networks only received visual sensory input. In contrast, our model received both visual and proprioceptive sensory inputs with realistic latencies, and the connections giving rise to the reaching behavior were not determined *a priori*. These differences allowed our model architecture to demonstrate sensory integration and the sensorimotor transformations that led to smooth, accurate reaching movements. Biologically inspired network models with visual and proprioceptive inputs, capable of adaptive voluntary reaching actions, tend to emphasize motor learning based on motor error, which has previously been correlated with cerebellar neural activity [9, 10]. Our model was not designed or tested for corrective reaching actions, nor attempting to model cerebellar cortex, however, evidence suggests that the feedback connections from PMd/M1 to PPC carry a corollary discharge signal that could lead to corrective reaches [2, 15]. Therefore, we would anticipate that the model could be evolved to perform online corrective reaching movements, however, the model could be missing critical components without a cerebellar region predicting the sensory consequences of motor commands (error signal).

Many movements, including reaches, show a rapid acceleration after movement onset followed by deceleration as the hand approaches the target. These smooth, accurate movements (Figure 3), which match control subject behavior for visually guided reaching [18, 19], emerged as a result of the evolved parameters. As for velocity profiles, single-humped profiles evolved, which indicates an optimal trade-off between speed and accuracy. In a multi-humped velocity profile, each hump represents a sub-movement towards the motor goal that indicates a non-optimal trade-off between speed and accuracy, and the magnitude of each sub-movement is dependent on the distance to the target and the target size [29]. We predict that the absence or manipulation of any parameters within the model should result in non-optimal corrective movements or a multi-humped velocity profile.

An interesting result came with the high accuracy of evolved agents. Theoretical evidence has indicated that sensory integration (from multiple modalities) leads to a better signal-to-noise ratio (SNR) provided that each sensory input is sufficiently noisy [30]. Based on this evidence, we predict that the integration of two distinct sensory inputs with different latencies provided sufficient sensory noise, leading to a better SNR, and resulting in behavioral accuracy greater than the implemented sensory resolution (Figure 4). Analysis of population diversity was done through MoI and average evolved parameter values (Table 1), and behavioral data (Figures 3-4). This analysis provided conclusive evidence that the ES found 100 independent solutions. As for the relationship between evolved parameters, the results suggest that a biological population should have some linear relationships between parameters governing neural firing patterns, and that topography would be preserved in feedforward but not in feedback projections. Finally, the results provide theoretical evidence for short-range excitation and long-range inhibition in recurrent connections across the PPC. Together these results provide a theoretical basis for future empirical studies investigating the fine-grained relationships between neural firing

patterns and the neural projections in a sensorimotor transformation task, such as reaching to a visual target.

The present study suggests how directly encoded evolved connectivity within an area such as the parietal cortex can lead to the construction of motor plans. These motor plans can overcome delays in sensory input, while still achieving accurate and rapid movements. Furthermore, predictions from the model suggest that there exists a diverse set of significant parameters governing these movements, and that manipulations of these parameters should lead to movement dysfunction. In future experiments, it would be of interest to implement indirect encoding of the weight parameters to investigate the importance of specific connectivity, and to see what types of excitation and inhibition emerge. In addition, it would be of interest to utilize the state of the art Covariance Matrix Adaptation (CMA) ES algorithm [31], and compare the evolutionary dynamics and evolved solutions with our current results.

5. REFERENCES

- [1] Mountcastle, V. B., Lynch, J. C., Georgopoulos, A., Sakata, H. and Acuna, C. Posterior Parietal Association Cortex of Monkey - Command Functions for Operations within Extrapersonal Space. *Journal of Neurophysiology*, 38, 4 (1975), 871-908.
- [2] Andersen, R. A. and Cui, H. Intention, action planning, and decision making in parietal-frontal circuits. *Neuron*, 63, 5 (Sep 10 2009), 568-583.
- [3] Sabes, P. N. Sensory integration for reaching, 1912011), 195-209.
- [4] Andersen, R. A. and Buneo, C. A. Intentional maps in posterior parietal cortex. *Annual Review of Neuroscience*, 25, 1 (2002), 189-220.
- [5] Desmurget, M., Epstein, C. M., Turner, R. S., Prablanc, C., Alexander, G. E. and Grafton, S. T. Role of the posterior parietal cortex in updating reaching movements to a visual target. *Nat Neurosci*, 2, 6 (Jun 1999), 563-567.
- [6] Caminiti, R., Chafee, M. V., Battaglia-Mayer, A., Averbeck, B. B., Crowe, D. A. and Georgopoulos, A. P. Understanding the parietal lobe syndrome from a neurophysiological and evolutionary perspective. *European Journal of Neuroscience*, 31, 12 (Jun 2010), 2320-2340.
- [7] Mulliken, G. H., Musallam, S. and Andersen, R. A. Inaugural Article: Forward estimation of movement state in posterior parietal cortex. *Proceedings of the National Academy of Sciences*, 105, 24 (2008), 8170-8177.
- [8] Cisek, P. Integrated Neural Processes for Defining Potential Actions and Deciding between Them: A Computational Model. *Journal of Neuroscience*, 26, 38 (2006), 9761-9770.
- [9] Passot, J.-B., Luque, N. R. and Arleo, A. Coupling internal cerebellar models enhances online adaptation and supports offline consolidation in sensorimotor tasks. *Frontiers in Computational Neuroscience*, 7(2013).
- [10] Shadmehr, R. and Krakauer, J. W. A computational neuroanatomy for motor control. *Exp Brain Res*, 185, 3 (Mar 2008), 359-381.
- [11] Xing, J. and Andersen, R. A. Models of the posterior parietal cortex which perform multimodal integration and represent space in several coordinate frames. *J Cogn Neurosci*, 12, 4 (Jul 2000), 601-614.
- [12] Zipser, D. and Andersen, R. A. A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, 331, 6158 (Feb 25 1988), 679-684.
- [13] Sober, S. J. and Sabes, P. N. Multisensory integration during motor planning. *J Neurosci*, 23, 18 (Aug 6 2003), 6982-6992.
- [14] Batista, A. P. and Andersen, R. A. The parietal reach region codes the next planned movement in a sequential reach task. *J Neurophysiol*, 85, 2 (Feb 2001), 539-544.
- [15] Buneo, C. A. and Andersen, R. A. The posterior parietal cortex: sensorimotor interface for the planning and online control of visually guided movements. *Neuropsychologia*, 44, 13 (2006), 2594-2606.
- [16] Mulliken, G. H., Musallam, S. and Andersen, R. A. Decoding trajectories from posterior parietal cortex ensembles. *J Neurosci*, 28, 48 (Nov 26 2008), 12913-12926.
- [17] Trappenberg, T. P. *Fundamentals of computational neuroscience*. Oxford University Press, Oxford ; New York, 2010.
- [18] d'Avella, A., Portone, A. and Lacquaniti, F. Superposition and modulation of muscle synergies for reaching in response to a change in target location. *J Neurophysiol*, 106, 6 (Dec 2011), 2796-2812.
- [19] Perfetti, B., Moissel, C., Landsness, E. C., Kvint, S., Pruski, A., Onofrij, M., Tononi, G. and Ghilardi, M. F. Temporal Evolution of Oscillatory Activity Predicts Performance in a Choice-Reaction Time Reaching Task. *Journal of Neurophysiology*, 105, 1 (2010), 18-27.
- [20] Keijzer, M., Merelo, J. J., Romero, G. and Schoenauer, M. Evolving objects: A general purpose evolutionary computation library. *Lect Notes Comput Sc*, 2310(2002), 231-242.
- [21] Kita, E. and Maruyama, T. *Genetic algorithm based on schemata theory*. City, 2011.
- [22] Morrison, R. W. and De Jong, K. A. Measurement of population diversity. *Lect Notes Comput Sc*, 2310(2002), 31-41.
- [23] Floreano, D., D'Unger, P. and Mattiussi, C. Neuroevolution: from architectures to learning. *Evolutionary Intelligence*, 1(2008), 47-62.
- [24] Stanley, K. O., D'Ambrosio, D. B. and Gauci, J. A hypercube-based encoding for evolving large-scale neural networks. *Artif. Life*, 15, 2 (2009), 185-212.
- [25] Van Geit, W., Achard, P. and De Schutter, E. Neurofitter: a parameter tuning package for a wide range of electrophysiological neuron models. *Front Neuroinform*, 1(2007), 1.
- [26] Rossant, C., Leijon, S., Magnusson, A. K. and Brette, R. Sensitivity of Noisy Neurons to Coincident Inputs. *Journal of Neuroscience*, 31, 47 (Nov 23 2011), 17193-17206.
- [27] Cecconi, F. and Parisi, D. Evolving organisms that can reach for objects. In *Proceedings of the Proceedings of the first international conference on simulation of adaptive behavior on From animals to animats* (Paris, France, 1990). MIT Press, [insert City of Publication],[insert 1990 of Publication].
- [28] Petrosino, G., Parisi, D. and Nolfi, S. Selective attention enables action selection: evidence from evolutionary robotics experiments. *Adaptive Behavior* 2013).
- [29] Meyer, D. E., Abrams, R. A., Kornblum, S., Wright, C. E. and Smith, J. E. Optimality in human motor performance: ideal control of rapid aimed movements. *Psychol Rev*, 95, 3 (Jul 1988), 340-370.
- [30] Sabes, P. N. Sensory integration for reaching: models of optimality in the context of behavior and the underlying neural circuits. *Prog Brain Res*, 191(2011), 195-209.
- [31] Igel, C., Hansen, N. and Roth, S. Covariance matrix adaptation for multi-objective optimization. *Evol Comput*, 15, 1 (Spr 2007), 1-28.