Neural reuse in multifunctional neural networks for control tasks

Lauren V. Benson[1,2], Madhavun Candadai[1,2] and Eduardo J. Izquierdo[1,2]

[1]Cognitive Science Program, Indiana University Bloomington
[2]The Luddy School of Informatics, Computing, and Engineering, Indiana University Bloomington
Corresponding email: lvbenson@iu.edu

Abstract

Living organisms perform multiple tasks, often using the same or shared neural networks. Such multifunctional neural networks are composed of neurons that contribute to different degrees in the different behaviors. In this work, we take a computational modeling approach to evaluate the extent to which neural resources are specialized or shared across different behaviors. To this end, we develop multifunctional feed-forward neural networks that are capable of performing three control tasks: inverted pendulum, cartpole balancing and single-legged walker. We then perform information lesions of individual neurons to determine their contribution to each task. Following that, we investigate the ability of two commonly used methods to estimate a neuron's contribution from its activity: neural variability and mutual information. Our study reveals the following: First, the same feed-forward neural network is capable of reusing its hidden layer neurons to perform multiple behaviors; second, information lesions reveal that the same behaviors are performed with different levels of reuse in different neural networks; and finally, mutual information is a better estimator of a neuron's contribution to a task than neural variability.

Introduction

As artificial intelligence, evolutionary robotics, and neuroscience become increasingly integrated, investigative efforts to understand the operation of neural networks is becoming increasingly important. An universal feature across living organisms is their ability to perform multiple behaviors. A predominant view of multifunctionality in neural networks involves utilizing distinct sub-networks within a larger neural network to perform the different behaviors (Dickinson, 1995; Tani et al., 2004; Bassett et al., 2011; Schrum and Miikkulainen, 2014). However, it has also been shown that the same dynamical neural network can be multifunctional in the presence of neuromodulation or plasticity (Chao et al., 2008; Sporns and Alexander, 2002; Yoder and Izquierdo, 2018) or even in its absence (Izquierdo and Bührmann, 2008; Williams and Beer, 2013; Agmon and Beer, 2014; Candadai and Izquierdo, 2018; Setzler and Izquierdo, 2017; Vasu and Izquierdo, 2017). In this work, we set out to artificially evolve neural networks for multiple control tasks without any assumptions of modularity or reuse and then investigate the degree to which neural resources are specialized or shared across tasks. Within evolutionary robotics, multifunctionality in neural networks has been primarily studied within the context of open-loop, input-output tasks (Yang et al., 2019a; Hong et al., 2020). Related work by Yang et al. (2019b) revealed functionally-specific clustering behavior when the same recurrent neural network is trained on various open-loop tasks. However, most tasks performed by living organisms are closed-loop behaviors involving a continuous interaction between the neural network and the environment. While efforts in evolutionary robotics have tackled learning to perform multiple closed-loop control tasks, to our knowledge artificial evolution has not been used to train non-modular multifunctional networks for more than two control tasks.

Our goals for this work are three-fold. The first aim is to extend previous efforts to evolve feed-forward neural networks to solve multiple closed-loop control tasks. Our second aim is to characterize the degree of neural reuse in the resulting successful multifunctional networks. We characterize neural reuse by estimating the contribution of each neuron to a task using information lesions on the hidden neurons of the neural network. However, in reality lesion studies may not be feasible and one must rely on analyzing the neural traces for insight into the agent's behavior. Our final aim is examine two methods that are commonly used to estimate a neuron's contribution to a task from neural activity alone: neural variability (Renart and Machens, 2014; Masquelier, 2013; Yang et al., 2019b), and mutual information (Wibral et al., 2017; Gabrié et al., 2018).

Methods

The goal of this paper is to train a neural network to perform multiple control tasks and then to analyze the resulting networks to study the extent to which neurons are reused across tasks. This section describes the tasks, the neural network model, the evolutionary optimization algorithm and finally the analysis methodologies used in this work.

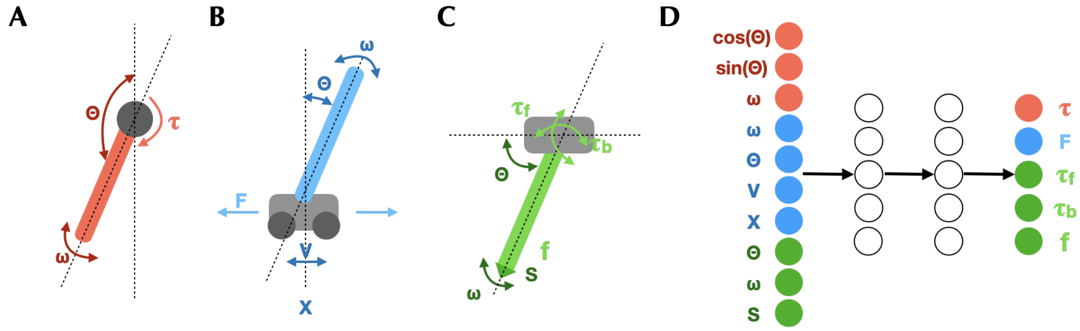

Figure 1: Task and neural network design. [A] Inverted pendulum task where the goal is to swing up the pendulum and keep in balanced. [B] Cartpole balancing task where the goal is to control the cart to keep the pole balanced [C] Single-legged walker task where the goal is control the foot swings and cover maximum distance in a fixed time. [D] Neural network architecture with distinct sensory and action units for each task and 2 hidden layers that are common across tasks.

Inverted pendulum The inverted pendulum swing-up task is a classical control task where the goal is to balance a pendulum by swinging it up with the least effort (Fig. 1A). Sensory observations for this task include $[cos(\theta), sin(\theta), \omega]$, where θ is the angle of the pendulum from the vertical and ω is its angular velocity. Controlling the pendulum involves applying a torque in the range $[-2, 2]$ at the center. Performance was evaluated by estimating the average cost over two trials each lasting 10 seconds, where is cost is estimated based on the applied torque. In each trial, the pendulum is initialized at the bottom ($\theta = \pi$) with an angular velocity of -1 and 1 respectively. For full specifications of the task in detail, see Brockman et al. (2016).

Cartpole balancing The cartpole balancing task is another classic control task that involves balancing a pole that is attached to a cart moving along a one-dimensional track. (Fig. 1B). Sensory observations for this task include the angle of the pole (θ), the pole's angular velocity (ω), the cart's position (x), and it's velocity ($\dot{x}$). The cart is controlled by applying a force to the cart in the range $[-1, 1]$, where the sign of the force determines the direction from which it is applied. Performance was evaluated as the average duration for which the pole was balanced over 4 trials that can each last up to 50 seconds. The pole was considered dropped if it falls beyond 12 degrees of the vertical, or if the cart moves more than 2.4 units from its initialized position. For full specifications of the task in detail, see Barto et al. (1983)

Single-legged walking The third control task implemented was the single-legged walker, where the goal is to control move a body along a one-dimensional rail by controlling a single-jointed leg (Fig. 1C). Sensory observations for this task include the angle of the leg (θ), the angular velocity (ω), and the binary foot state (the foot is either up or down at any given time). The leg is controlled by applying forces on its forward-backward swing, its up-down swing and by putting the foot up or down. The walker's body changes position only when the foot is in the down state. The walker is initialized with its leg geometry in the forward position each time. Performance is evaluated as the average distance covered over 9 trials that each lasted 110 seconds. For full specifications of the task in detail, see Beer and Gallagher (1999).

Neural network As the neural controller of the agent, we used a feed-forward neural network with two hidden layers containing 5 neurons. The input layer had 10 units corresponding to the total number of sensory inputs of all tasks and similarly the output layer had 4 units corresponding to the actions of each task (Fig. 1D). The hidden layer activations were set to be ReLU, $f(x) = max(0, x)$ and output layer activations were sigmoidal, $f(x) = 1/(1 + exp(-x))$.

Evolutionary algorithm Parameters of the neural network model were optimized using a real-valued microbial genetic algorithm Harvey (2009). The trend in machine learning has been to create increasingly sophisticated optimization algorithms and to attribute increases in performance to the bells and whistles. For this work we employed a deliberately minimal version of an artificial evolutionary algorithm to demonstrate selection, inheritance, and variation are sufficient to generate neural networks capable of multifunctional behavior. The parameters optimized included the weights between all layers and the biases of each neuron, totaling 115 parameters. We first evolve neural networks for a single task, where the fitness is given by the performance of the neural network in that task. We then evolve a neural network to solve all three tasks. In this case, the fitness of the individual was given by the product of fitness in each task. This guarantees good performance across all tasks. Each evolutionary search was initialized with 50 random individuals and evolved over 50000 tournaments.

Information lesions Information lesions have appeared in the neuroscience literature in a variety of a context to study particular brain functions (Vaidya et al., 2019; Koenigs et al., 2007). In our work, we use information lesions to determine the role that each neuron plays in each task for any one multifunctional neural network. We use this as the 'ground truth' of neural reuse for that neural network. The idea is to systematically disable each neuron in the hidden layer and study its consequence on behavioral performance. Using ablations to cut off all outgoing information from each neuron has too harsh of an effect on the system, creating opportunities for unjustified neural activity. Instead, we created information lesions by fixing the neuron's outputs to a constant value. This value is chosen by performing a sweep across the entire range of that neuron's activations to determine the value that has the most impact on behavioral performance.

Neural variability A systematic information lesion analysis is not always possible in an experimental setting. In this work, we examine less intrusive predictive measures that rely on neural activity alone. As a first predictive measure for whether a neuron contributes to performing a behavior, we calculated the variance in a neuron's activity during the course of a task. Intuitively, higher variance in activity is indicative of greater involvement in the task. This is given by the variance of the state of the neuron for that task, $1/n \sum_{i=1}^{n}(x_i - \mu)^2$ where μ refers to the mean neural activity of the given neuron, x_i refers to an i^{th} sample of neural activity and n refers to the total number of samples.

Mutual information As a second predictive measure of a neuron's role in a behavior, we determined the mutual information between the sensory input and a neuron's activity. Again, it is intuited that the mutual information is directly proportional to the contribution of a neuron in a task. Mutual information (MI) is given by:

$$MI = \sum_{y \in Y} \sum_{x \in X} p(x, y) \; log \frac{p(x, y)}{p(x)p(y)} \tag{1}$$

where X and Y correspond to the sensory input and neural activity respectively, $p(x, y)$ refers to their joint probability and $p(.)$ their marginal probabilities. MI was estimated using the *infotheory* package (Candadai and Izquierdo, 2019) in which probability densities were estimated using an average-shifted histogram approach (Scott, 1985).

Results

To what degree are multifunctional neural networks either reusing neural resources or dedicating specialized circuitry for each task? In order to address this question, we first set out to artificially evolve an ensemble of neural networks on multiple control tasks. We then characterized the degree of reuse and specialization in the neurons of the multifunctional networks using informational lesions. Finally, we evaluated two statistical measures on their abilities to estimate function from neural traces.

Evolving neural networks for each individual task

In order to evolve a neural network capable of solving multiple tasks, we first verified that each of the tasks could be evolved using a similar neural network and fitness functions. This first experiment allowed us to ensure that: (a) the size and architecture of the neural network is appropriate for solving each task; (b) the task-specific parameters and constraints of the sensory inputs are likewise appropriate for solving each task; and (c) the fitness function used for each task reliably results in successfully evolved neural networks. For each task, we performed 10 independent evolutionary runs with different random seeds for 150 generations. In preliminary analysis, we noted that neural networks that performed the task successfully had a fitness of 0.93 or higher. The fitness for each task was normalized to run between 0 and 1. For the inverted pendulum, 9 out of the 10 evolutionary runs produced neural networks that could solve the task with a fitness greater than the 0.93 threshold. The best neural network for this task obtained a fitness of 0.96. All 10 of the runs for the cartpole balancing task resulted in successful neural networks, with the best neural network obtaining a fitness of 0.99. Of the three, the single-legged walking task was the hardest to solve, with only four of the 10 runs producing neural networks with a fitness greater than 0.93. However, the best neural network for this task also obtained a fitness of 0.99. Given the success obtained with the size and architecture of the neural network, the arrangement of the sensory input and motor output, and the shape of the fitness evaluations, it is in principle possible that a single neural network can solve all three tasks.

Evolving neural networks for multiple tasks

The majority of work training artificial neural networks to solve tasks has focused on single tasks, however biologi-

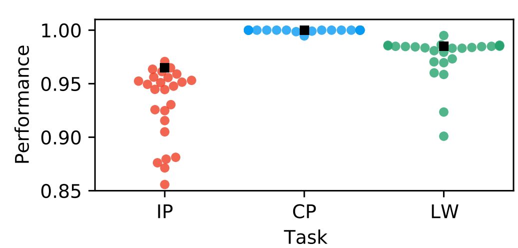

Figure 2: Performance of the filtered ensemble on individual tasks. Each point represents the performance on a fine-grained set of starting conditions of each circuit in the ensemble, for each of the tasks. The neural network with the highest overall fitness is shown in black. All neural networks in this set are successful multifunctional neural networks.

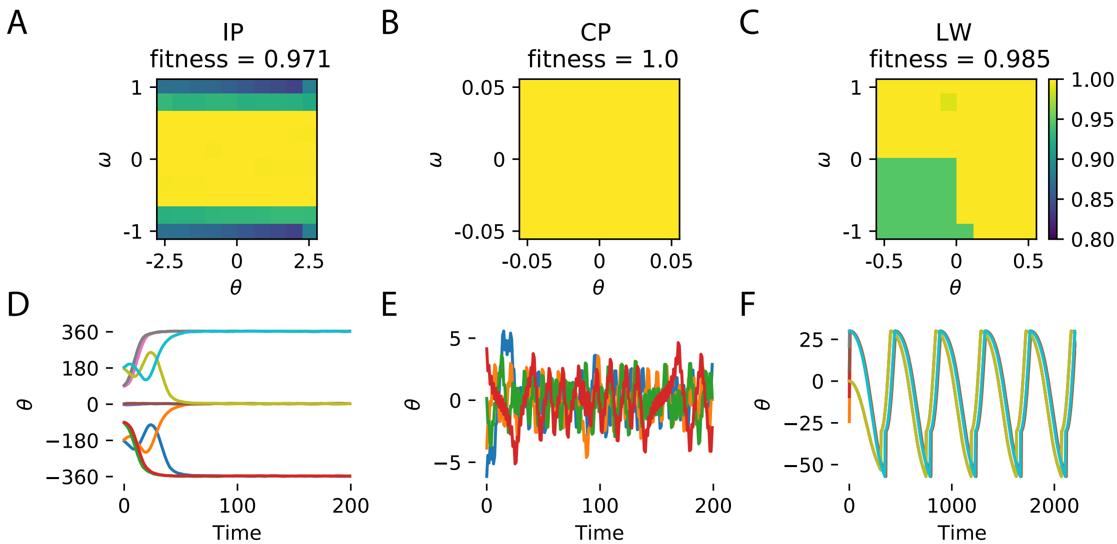

Figure 3: Behavior of the best performing neural network. Top row: Performance maps as a function of the starting conditions for each task. Bottom row: Behavioral traces for a sample of starting conditions for each task. [A] Performance as a function of the initial angular velocity and angle of the pendulum in the inverted pendulum task. The color represents performance. Yellow represents perfect performance and dark blue represents 80% performance. [B] Performance as a function of the initial angular velocity and angle of the pole in the cartpole balancing task showing perfect performance throughout. The position and velocity of the cart was also varied systematically, with the performance averaged across those two dimensions. [C] Performance as a function of the initial angular velocity and angle of the leg in the single-legged walker. Yellow represents perfect performance and green represents 95% performance. [D] Sample trials from the inverted pendulum task showing the best neural network swinging the pendulum up to 0 (equivalent to 360 and -360). [E] Sample trials from the cartpole balancing task showing the best neural network oscillating the pole and maintaining it near the vertical. [F] Sample trials from the single-legged walker showing leg angle over time, indicating successful walking behavior by the best neural network.

cal neural networks can produce multiple different behaviors seamlessly. The next step in this work was to evolve the same neural network to solve all three tasks. The motivation for this step is two-fold. First, we would like to further develop a methodology to generate multifunctional neural networks for control tasks. Second, we will use the generated multifunctional neural networks to investigate neural reuse, and the tools of analysis that allow us study neural reuse. All evolutionary runs achieved some success on the three tasks. Out of these 100 runs, 32 neural networks solved the three tasks with a combined fitness greater than 0.80, within a fitness domain of [0,1]. We selected the ensemble of neural networks evolved in those runs to analyze in more detail.

Ensemble performance on individual tasks

The fitness of these neural networks was calculated as the product of their performance across the three tasks on a relatively small subset of starting conditions (27 starting conditions in total). To ensure that the neural networks in the ensemble could indeed solve each of the tasks robustly, we analyzed their performance on a finer grained set of starting conditions for each of the tasks. Specifically, we examined the performance across 10 different starting angles for the pendulum/pole/leg, and across 10 different starting angular velocities (300 conditions total). Next, we filtered the ensemble to those that had a performance of 0.85 or higher on each of the individual tasks on the finer-grained analysis. This resulted in a total of 25 neural networks (Fig. 2). Given their high level of performance across a wide range of starting conditions, consistently for each of the tasks, the neural networks in this filtered ensemble can be used for our analysis of neural reuse. The neural network with the highest overall fitness had a performance of 0.97 on the inverted pendulum task, 1.0 on the cartpole balancing task, and 0.98 on the single-legged walking task.

Behavior of the best performing neural network

In order to further validate the meaning of these performance scores across the three tasks, we visualized the robustness across starting conditions and behavior for one neural network (Fig. 2), the best performing one from the ensemble. We examined the robustness of the neural network for the inverted pendulum (Fig. 2A), the cartpole balancing task (Fig. 2B), and the single-legged walker (Fig. 2C) across a

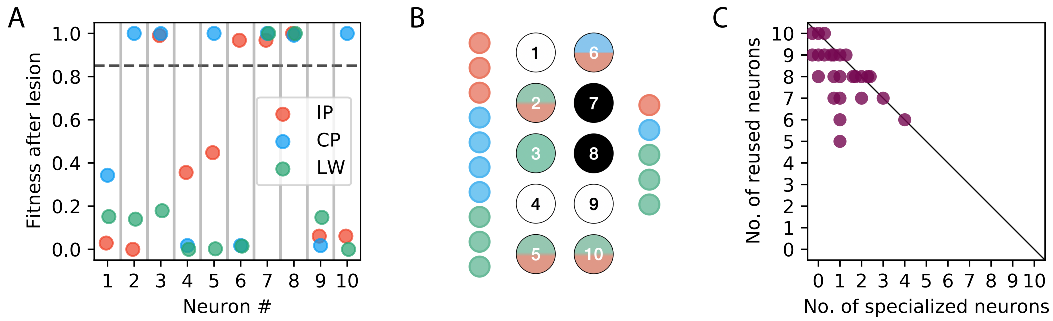

Figure 4: Neural reuse in the best and the ensemble of multifunctional neural networks. [A] Task performance after information lesions were performed on each neuron in each task. For each neuron (horizontal axis), orange represents performance in inverted pendulum after lesioning that neuron, blue in cartpole balancing and green in single-legged walker. A neuron is said to significantly contribute to a task if performance drops below a chosen threshold of 0.85 upon information lesioning of that neuron. [B] Neural network representation (same architecture as in figure 1D) of each neuron's contribution to the three tasks. Tasks are represented by the same colors as in Panel A. White denotes contribution to all tasks and black denotes contribution to none. When a neuron contributes to more than one task it is colored using both colors. Of the 8 hidden layer neurons that are contributing to at least one task 7 are reused, in the best neural network. [C] Number of specialized neurons contributing to one of the three tasks versus number of reused neurons contributing to at least two of the three tasks. Each dot represents one of the 25 filtered neural networks. All neural networks show high levels of reuse rather than specialization.

wide range of starting angles for the pole/leg and starting angular velocities. Overall, the performance of the neural network is consistently good across the full range of starting conditions for all three tasks. We also examined the behavior of this neural network over time, for a subset of starting conditions for the three tasks (Fig. 2D-F). We specifically examined the starting angles for the pole/leg in each task, and, as expected, the behavior of the neural network in each of the tasks is consistent with its performance: the inverted pendulum gets balanced regularly (Fig. 2D), the pole is maintained near the top of the cart (Fig. 2E), and the leg is moved back and forth, while the foot is lifted and lowered to produce forward movement (Fig. 2F).

Neural reuse in multifunctional neural networks

The main contribution of this paper is to better understand how neural networks are multifunctional. The work up to this point allowed us to arrive at an ensemble of neural networks that can solve multiple tasks to a high degree of competency. The main advantage of this is that we now have complete access to these multifunctional neural networks to further examine how they accomplish this feat. The goal of this section is to characterize which neurons in the network contribute to which tasks and to study the variance observed across different neural networks in the ensemble. In order to do this, we systematically lesion each interneuron in the neural networks and measure its effect on performance for each task (Fig. 4). We do this analysis first for the best performing neural network (Fig. 4A and B), and then we generalize the analysis for the rest of the neural networks in the ensemble (Fig. 4C). Instead of a traditional lesion, where the target neuron is fixed to zero, we performed an information lesion, effectively allowing for the target neuron to take the fixed value that maximized performance. Although more computationally costly, an information lesion allows us to more finely dissect the functional role of each neuron in the network.

In order to understand the effect of information lesions on the different tasks, we first analyzed the results from the analysis on the best performing neural network (Fig. 4A). As expected, each neuron participates in each of the tasks to different degrees. Interestingly, the performance disruption from the information lesions is relatively binary: a lesion to any one specific neuron either disrupts the performance for a task gravely (for example the performance drops more than 50% of its usual level), or it does not affect it much (for example the neural network remains with a performance above 80%). This suggests a natural range to define a threshold. For the purpose of this analysis, we consider a neuron involved in a task if the information lesion disrupts performance below 85%. When we take all three tasks into consideration, each neuron can be categorized into one of seven categories: the neuron does not contribute to any task, it contributes to only one of the three tasks, it contributes to only two of the tasks, or it contributes to all three tasks. In the case of the best performing neural network (Fig. 4B), we can see that two neurons (7 and 8) do not contribute to any task; only one neuron (3) is dedicated to a single task; four neurons are dedicated to two tasks (neurons 3, 5, and 8 are dedicated to the inverted pendulum and single-legged walker and

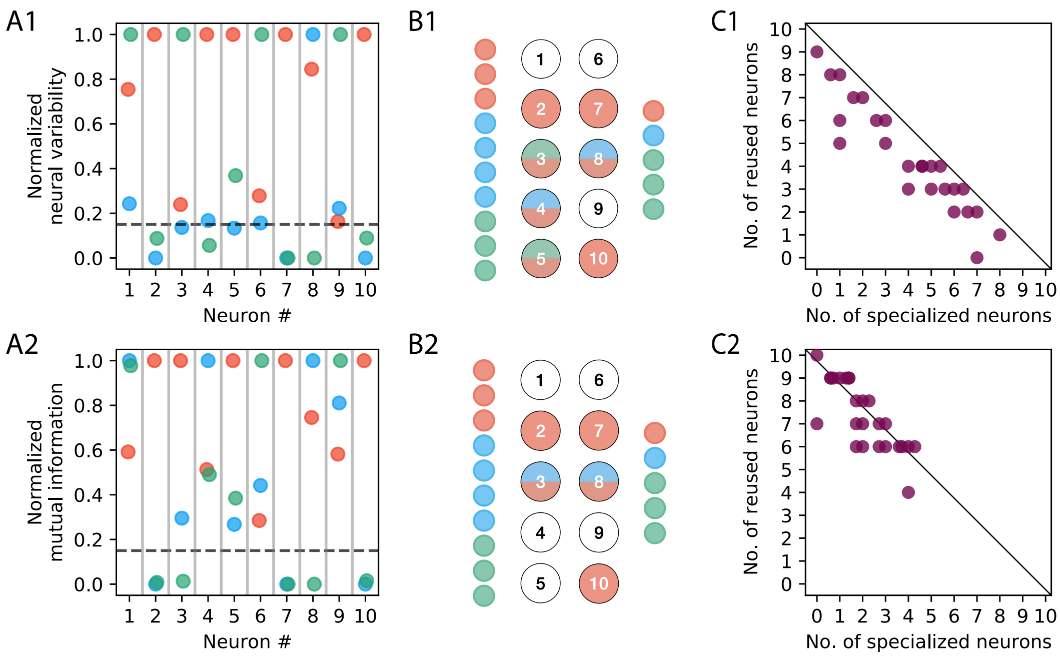

Figure 5: Estimating neural reuse from neural activity. Top row: Analysis of best and ensemble using neural variability. Bottom row: Analysis of best and ensemble using mutual information [A1] Normalized neural variability in each neuron estimated from its activity during each task. As opposed to information lesions, higher values correspond to greater contribution. Hence, a neuron is said to contribute to a task if it has neural variability over 0.15. [B1] Neural network representation (same architecture as in figure 1D) of each neuron's contribution to the three tasks. Tasks are represented by colors same as in panel A. White denotes contribution to all tasks and black denotes contribution to none. When a neuron contributes to more than one task it is colored using both colors. According to neural variability, 7 out of 10 hidden layer neurons in the best neural network are reused. [C1] Number of specialized neurons contributing to one of the three tasks versus number of reused neurons contributing to at least two of the three tasks. Each dot represents one of the 25 filtered neural networks. Neural variability is relatively more rigid in designating reuse. [A2] Same as A1 but for mutual information. [B2] Same as B2 but for mutual information. According to mutual information, 7 out of 10 hidden layer neurons in the best neural network are reused. [C2] Same as C1 but for mutual information. Mutual information shows a higher level of reuse compared to neural variability.

neuron 6 is dedicated to the cartpole balancing and Single-legged walker); and three neurons (1, 4, and 9) are dedicated to all three tasks. The main point of interest with this analysis is that most of the neurons are highly multi-functional: the neurons are reused across multiple tasks.

Is the degree of neural reuse observed in the best performing multifunctional neural network similar for other neural networks in the ensemble? In order to examine the degree of reuse and specialization across the full ensemble of successful neural networks, we further simplified the categories of each neuron to either not involved in any task or if involved, then either specialized or reused, depending on whether the neuron was involved in one or multiple tasks, respectively. Interestingly, we found a relatively consistent pattern of neural reuse that evolved across the multifunctional neural networks in the ensemble (Fig. 4C): the majority of neurons in nearly all of the neural networks were being reused across two or all three tasks.

Estimating neural reuse from neural activity

Although we have demonstrated that performing an exhaustive lesion study allows us to determine the degree of neural reuse in an artificial system, such a level of manipulation would not be practical for most living organisms. In this last section, we analyze the degree to which we can estimate neural reuse using only neural traces generated from ongoing behavioral recordings. Specifically, we examine two of the primary methods used to analyze neural traces: neural variance and mutual information between each neuron and the sensory stimuli. We first examined the neural variability (Fig. 5A1) and mutual information (Fig. 5A2) for the best performing neural network across the three tasks. Note that in the informational lesions studies, low values represented likely involvement of that neuron in that task, whereas in the neural variance and mutual information analysis the opposite is true: high values represent likely involvement in a task. Following the analysis done for the lesions, we categorized each of the neurons as contributing to different subsets of the tasks (Figs. 5B1 and 5B2), respectively for each metric. We used the same threshold (85%) as in the lesion studies to classify neuron involvement. Note that 85% here is represented as 15% given the inverse logic of the data. The analysis of the best neural network reveals a similar pattern for both neural variance and mutual information. Neural variance estimates perfectly the role of 3 out of the 10 neurons (1, 5, and 9), and partially the role of 5 neurons (2, 3, 4, 6, and 10). Mutual information also estimates perfectly the role of 3 neurons (1, 4, 9) and partially the role of 5 (2, 4, 5, 6, and 10). Finally, in order to examine the estimates of the degree of reuse and specialization across the full ensemble of successful neural networks, we again simplified the categories of each neuron to either specialized or reused (Figs. 5C1 and 5C2, respectively for each metric). Qualitatively, the neural variance analysis suggests there is a relatively balanced set of neural networks in the ensemble across the specialized and reused spectrum; whereas the mutual information analysis suggests the neural networks are mostly comprised of neurons with multifunctional roles. From the lesion analysis, we know the latter is closer to the ground truth, as visualized by the comparison of Figs. 4C, 5C1 and 5C2. Finally, in order to assess the relative merit of the two metrics quantitatively, we measured the Euclidean distance between each of the two metrics (neural variance and mutual information) against the information lesion, across all 25 neural networks in the ensemble, using the raw data in Figures 4A, 5A1, and 5A2 (see Fig. 6A), as well as for the categorical data in Figures 4C, 5C1, and 5C2 (see Fig. 6B). For both the raw data (Fig. 6A) and the categorical data (Fig. 6B), the the mutual information consistently estimated the information lesion data more closely.

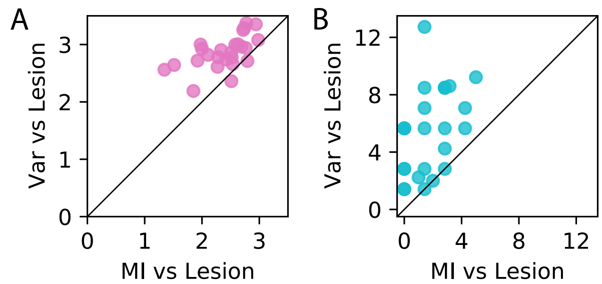

Figure 6: Comparing metrics for estimating neural reuse. Euclidean distance between the normalized metrics (neural variance and mutual information) and the information lesions for both the raw and categorical data. [A] Euclidean distance of normalized mutual information in each neuron for each task, and normalized variance of each neuron for each task (from the normalized fitness values after lesioning each neuron) plotted against one another. Each dot is a neural network from the ensemble. Most dots lie above the diagonal denoting that mutual information estimates are closer in euclidean distance to lesion scores than neural variability. [B] Euclidean distance of the number of specialized and reused neurons given by mutual information and neural variability to that given by the information lesions plotted against each other. Each dot is a neural network from the ensemble. Like in panel A, mutual information estimates are closer in euclidean distance to lesion scores than neural variability.

Discussion

Our work extends previous studies in understanding multifunctionality by developing a neural network that can perform multiple control tasks and by investigating its extent of neural reuse. Our results show that (1) evolutionary algorithms can be employed to successfully evolve feed-forward neural networks to solve multiple control tasks, (2) using information lesions to study neural reuse in our multifunctional neural networks reveals that those with similar behavioral performance nevertheless differ in the extent of specialization and reuse, and (3) mutual information outperforms neural variability as a method to evaluate neural reuse using only the neural activity.

Our work has both practical and theoretical implications. From a practical perspective, the evolutionary algorithm used in our work demonstrates that this rather straightforward alternative to several contemporary algorithms in artificial intelligence (Ruder, 2017) can generate multifunctional neural networks. Furthermore, evaluating performance on all tasks at every generation prevents potential issues that may arise with the use of reinforcement learning and backpropagation, such as catastrophic interference. From a theoretical standpoint, the information lesions performed in our analysis is a lesioning approach that is an unbiased alternative to ablating the neuron altogether. Furthermore, our preliminary investigations provide provide insights into the relative abilities of other statistical methods to infer the ground truth obtained from the lesion studies.

The work presented in this paper provides ample opportunities for future expansion. In our analysis of a neuron's contribution to each task, the information lesion analysis was considered the "ground truth." In experimental settings, such a systematic set of lesions across all components of a system is challenging and often not feasible. For this reason, we examined two measures that relied only on the system's activity over time while performing the multiple behaviors. Specifically, we used neural variability and mutual information as estimators of a neuron's role in a behavior from neural activity alone. A limitation of the current work is that we only considered the role of each individual neuron. This assumes that a neuron's contribution to a behavior is independent of other neurons in the neural network. In future work, we intend to inspect these neural networks over different combinations of grouped neurons, rather than exclusively on an individual level. This may reveal additional insight into redundant or synergistic reuse of neurons in individual tasks as well as across tasks. We also plan to inspect our neural reuse classifications over a wider range of thresholds. Furthermore, in this paper, we analyzed patterns of neural reuse in a feed-forward neural network of a particular size and architecture, and we would like to expand this analysis to neural networks of different sizes and architectures. To what extent do the patterns that we have identified suggest similar patterns in neural networks with larger or smaller hidden layers? We are also interested in continuing to investigate the multifunctionality of these neural networks by adding more closed-loop control tasks to the framework, and the resulting neural reuse patterns. We would also like to continue to use neural traces to predict the levels of neural reuse identified from information lesions with additional statistical methods. Finally, we intend to further harness the sophisticated behavioral capabilities of biological systems by expanding our neural model to include continuous-time dynamical recurrent neural networks. Developing these synergies between artificial models and living organisms paves the way for designing increasingly realistic and behaviorally robust artificial systems.

Data availability

The simulation code and data files are publicly available in our research group's GitHub account: github.iu.edu/EASy/BensonALife2020.

Acknowledgements

We are thankful for the reviewers' comments. LVB is supported by NSF-NRT grant 1735095, Interdisciplinary Training in Complex Networks and Systems. This material is based upon work supported by the National Science Foundation under Grant No. 1845322.

References

Agmon, E. and Beer, R. D. (2014). The evolution and analysis of action switching in embodied agents. *Adaptive Behavior*, 22(1):3–20.

Barto, A. G., Sutton, R. S., and Anderson, C. W. (1983). Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE transactions on systems, man, and cybernetics*, (5):834–846.

Bassett, D. S., Wymbs, N. F., Porter, M. A., Mucha, P. J., Carlson, J. M., and Grafton, S. T. (2011). Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences*, 108(18):7641–7646.

Beer, R. and Gallagher, J. (1999). Evolution and analysis of model cpgs for walking: Ii. general principles and individual variability. *Computational Neuroscience*, (7):119–147.

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. (2016). Openai gym. *CoRR*, abs/1606.01540.

Candadai, M. and Izquierdo, E. (2018). Multifunctionality in embodied agents: Three levels of neural reuse. *arXiv preprint arXiv:1802.03891*.

Candadai, M. and Izquierdo, E. J. (2019). infotheory: A c++/python package for multivariate information theoretic analysis. *arXiv preprint arXiv:1907.02339*.

Chao, Z. C., Bakkum, D. J., and Potter, S. M. (2008). Shaping embodied neural networks for adaptive goal-directed behavior. *PLoS Computational Biology*, 4(3).

Dickinson, P. S. (1995). Interactions among neural networks for behavior. *Current Opinion in Neurobiology*, 5(6):792 – 798.

Gabrié, M., Manoel, A., Luneau, C., Macris, N., Krzakala, F., Zdeborová, L., et al. (2018). Entropy and mutual information in models of deep neural networks. In *Advances in Neural Information Processing Systems*, pages 1821 – 1831.

Harvey, I. (2009). The microbial genetic algorithm. In *European conference on artificial life*, pages 126–133. Springer.

Hong, C., Wei, X., Wang, J., Deng, B., Yu, H., and Che, Y. (2020). Training spiking neural networks for cognitive tasks: A versatile framework compatible with various temporal codes. *IEEE Transactions on Neural Networks and Learning Systems*, 31(4):1285–1296.

Izquierdo, E. and Bührmann, T. (2008). Analysis of a dynamical recurrent neural network evolved for two qualitatively different tasks: Walking and chemotaxis. *ALIFE.*, pages 257–264.

Koenigs, M., Tranel, D. T., and Damasio, A. R. (2007). The lesion method in cognitive neuroscience. *Handbook of Psychophysiology*, 3:139–156.

Masquelier, T. (2013). Neural variability, or lack thereof. *Frontiers in Computational Neuroscience*, 7:7.

Renart, A. and Machens, C. (2014). Variability in neural activity and behavior. *Current opinion in neurobiology*, 25C:211–220.

Ruder, S. (2017). An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.

Schrum, J. and Miikkulainen, R. (2014). Evolving multimodal behavior with modular neural networks in ms. pac-man. In *Proceedings of the 2014 annual conference on genetic and evolutionary computation*, pages 325–332.

Scott, D. W. (1985). Averaged shifted histograms: Effective nonparametric density estimators in several dimensions. *Ann. Statist.*, 13(3):1024–1040.

Setzler, M. and Izquierdo, E. J. (2017). Adaptability and neural reuse in minimally cognitive agents. In *CogSci*.

Sporns, O. and Alexander, W. H. (2002). Neuromodulation and plasticity in an autonomous robot. *Neural Networks*, 15(4):761 – 774.

Tani, J., Ito, M., and Sugita, Y. (2004). Self-organization of distributedly represented multiple behavior schemata in a mirror system: reviews of robot experiments using rnnpb. *Neural Networks*, 17(8-9):1273–1289.

Vaidya, A., Pujara, M., Petrides, M., Murray, E., and Fellows, L. (2019). Lesion studies in contemporary neuroscience. *Trends in Cognitive Sciences*, 23.

Vasu, M. C. and Izquierdo, E. J. (2017). Evolution and analysis of embodied spiking neural networks reveals task-specific clusters of effective networks. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 75–82.

Wibral, M., Priesemann, V., Kay, J. W., Lizier, J. T., and Phillips, W. A. (2017). Partial information decomposition as a unified approach to the specification of neural goal functions. *Brain and Cognition*, 112:25 – 38.

Williams, P. and Beer, R. (2013). Environmental feedback drives multiple behaviors from the same neural circuit. *Artificial Life Conference Proceedings*, (25):268–275.

Yang, G., Cole, M., and Rajan, K. (2019a). How to study the neural mechanisms of multiple tasks. *Current Opinion in Behavioral Sciences*, 29:134–143.

Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T., and Wang, X.-J. (2019b). Task representations in neural networks trained to perform many cognitive tasks. *Nature neuroscience*, 22(2):297–306.

Yoder, J. A. and Izquierdo, E. J. (2018). Behavioral stability in the face of neuromodulation in brain-body-environment systems. In *Artificial Life Conference Proceedings*, pages 268–275. MIT Press.