Analysis Pipeline for High-Dimensional Neuromechanical Model Improvement

Camila J. Fernandez[1(✉)], Jeffrey M. McManus[4], Yanjun Li[5], Michael J. Bennington[1], Roger D. Quinn[5], Hillel J. Chiel[4,6,7], and Victoria A. Webster-Wood[1,2,3]

[1] Department of Mechanical, Carnegie Mellon University, Pittsburgh, PA 15232, USA
{cjfernan,vwebster}@andrew.cmu.edu
[2] Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA 15232, USA
[3] McGowan Institute for Regenerative Medicine, Carnegie Mellon University, Pittsburgh, PA 15232, USA
[4] Department of Biology, Case Western Reserve University, 10900 Euclid Ave, Cleveland, OH 44106, USA
[5] Mechanical Engineering, Case Western Reserve University, 10900 Euclid Ave, Cleveland, OH 44106, USA
[6] Neurosciences, Case Western Reserve University, 10900 Euclid Ave, Cleveland, OH 44106, USA
[7] Biomedical Engineering, Case Western Reserve University, 10900 Euclid Ave, Cleveland, OH 44106, USA

Abstract. To capture and understand animal behavior, engineers and biologists seek to develop biologically accurate neuromechanical models of muscle dynamics and neural control. However, demand-driven enhancement of complex neuromechanics, such as the multifunctional *Aplysia californica* feeding apparatus, can be challenging due to the multidimensional biomechanical and neural models involved. We propose an analysis pipeline that enables reinforcement learning (RL) to classify which aspects of an engineered neuromechanical model can accurately capture animal behavior. As an example, prioritizing where demand-driven enhancement of a biomechanical and neural model is needed, the neural model of a published neuromechanical model of *Aplysia* swallowing during feeding was replaced with an RL controller and their performances were compared and correlated with *in vivo* swallowing behavior. By comparing the performance of the neural model and the learned model to *in vivo* animal behavior, we can pinpoint areas for improvement. The analysis pipeline identified that the neuromechanical model confidently captured force performance with no significant difference from

This work was supported by NSF DBI2015317 as part of the NSF/CIHR/DFG/FRQ/UKRI-MRC Next Generation Networks for Neuroscience Program, by the NSF Research Fellowship Program under Grant No. DGE1745016, and by internal funding through Carnegie Mellon University. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

N. S. Szczecinski et al. (Eds.): Living Machines 2024, LNAI 14930, pp. 333–348, 2025.
https://doi.org/10.1007/978-3-031-72597-5_23

animal swallowing force behavior. It most usefully also indicated that the biomechanical model will need to be improved in future iterations to better capture motor neuron activity. Future work should explore the accuracy of the RL-enabled analysis pipeline with a more advanced biomechanical model.

Keywords: reinforcement learning · neuromechanical model · *Aplysia californica*

1 Introduction

Biologically-accurate neuromechanical models of complex animal behavior are valuable tools for engineers and biologists. Engineers and roboticists often use these models to learn how to design robots capable of precise control and robust adaptability witnessed in nature [1–4]. Biologists and neuroscientists can test biological hypotheses using these models to deepen their understanding of what has been observed experimentally and in literature [5–8]. One common model organism for developing these neuromechanical models is the marine mollusk *Aplysia californica*. This animal displays complex multifunctionality in its feeding behaviors despite its relatively small neuromusculature. With this tractable neuromusculature system, *Aplysia* can adjust its feeding behavior between biting, swallowing, and rejecting food (i.e., seaweed) efficiently [9–12]. For example, if *Aplysia* detects that an object it was ingesting is inedible, it can switch from swallowing behavior to rejection behavior. By developing models that represent muscle dynamics and neural control, roboticists and biologists alike can better understand how *Aplysia* is capable of this multifunctionality and how this could be applied to advancing current soft robots [13–15].

When developing models that can represent the multifunctional behavior of *Aplysia's* feeding apparatus, the question of how complex the model should be becomes very important due to the trade-off between low computational cost and high complexity. Therefore, model complexity should be dependent on the application of the model. For example, in robotics, one may want a fast model that is moderately biologically accurate, whereas in neuroscience, one may want a highly accurate model, even if it is computationally slower. Given this trade-off, a model should be as complex as needed for a goal, but no more so. This compromise is at the core of demand-driven complexity [13]. However, in complex neuromechanical models, it can be challenging to identify what aspects of the model need to be refined at any given stage of model development.

In this work, we present a method that leverages reinforcement learning (RL) in an analysis pipeline to identify what aspects of a neuromechanical model of *Aplysia* need to be improved towards better capturing *in vivo* animal behavior. Our approach both identifies which components in the tested model deviate from *in vivo* animal behavior and gives a potential order of priority for which components need improvement.

2 Methods

2.1 Analysis Pipeline

Systematically identifying what aspects of a model to focus on when using demand-driven complexity in neuromechanical modeling is challenging because there are currently few quantitative tools to assess the aspects of these highly complex models that deviate most from the target data. For the approach presented here, we focused on identifying if the biomechanical model, neural model, or both should have priority for future model improvements. Our analysis pipeline begins with the GymSlug Reinforcement Learning Gym [16]. GymSlug is based on a previously published neuromechanical model of *Aplysia* feeding [13] in which the neural circuitry is represented using Boolean operations, and the biomechanics are a simplified quasistatic spring model with 1^{st} order muscle dynamics. The model parameters were hand-tuned simultaneously by experts in *Aplysia* feeding behavior to produce qualitatively similar multifunctional feeding behavior. In GymSlug, the neural model is replaced with an RL controller, which takes the state of the biomechanics as inputs and outputs motor neuron activities. We posit that if the biomechanical model were a perfect representation of the animal biomechanics, given enough training data, then the RL model would converge to a control strategy that highly correlated with the *in vivo* neural signals. We will call the expert hand-tuned neuromechanical model the "Expert" model and the RL policy model the "Trained" model.

To identify areas within the Expert model that may need demand-driven complexity-based improvements, we first ran the GymSlug training program on six randomly-seeded iterations with a training period of 500,000 epoch iterations to match the amount of experimental *in vivo Aplysia* data. This created six individual RL control policies, which were used to generate simulated data for a sequence of four to five swallows each following training.

Following model data generation, all the experimental and predicted animal swallowing data was segmented and normalized. We then calculated the cross-correlation coefficient distribution of fundamental motor and interneural swallowing behavior properties between the Animal data, Expert models, and Trained models. Inter-animal, Animal-to-Expert, and Animal-to-Trained model cross-correlation distributions were compared to determine if the predicted swallowing behavior is highly correlated with *in vivo* swallowing behavior. Using this approach, if the Expert model did not highly correlate with animal data and the learned neural activations of the Trained model resulted in a higher correlation with the animal data than the neural model, improvement for those components in the neural model should be a high priority. If the learned activations are not more highly correlated, then improvements to the biomechanical model should be a priority.

Animal Swallowing Dataset. To test the ability of our analysis pipeline to quantitatively identify discrepancies between animal behavior and neuromechanical models, a diverse dataset of animal behavior and model-predicted behavior

is required. The animal data was collected in [9]. Briefly, adult *Aplysia* were placed in a temperature-controlled testing chamber at 14–16°C, and a piece of unbreakable seaweed, composed of two strips of Nori applied on either side of double-sided tape, was hung over this chamber attached to a load cell. Electrodes were implanted in the *Aplysia* to record activity in the I2 muscle, the radular nerve, buccal nerve 2, and buccal nerve 3. For each experiment, 4 to 5 swallow cycles were extracted from the raw data since that is the observed favorable time until the slug begins to behave differently [9]. The dataset from this work consists of six feeding sequences with four to five consecutive swallows from five individual *Aplysia* [9]. Unbreakable seaweed was used to allow swallowing pulling force measurement and recording activity in key motor neurons. This pipeline includes the activation frequency of the I2 muscle as well as the firing frequencies of B8a/b (manipulating the grasper closing I4 muscle), B38 (activating the pinching function in I3), and B6/B9/B3 (controlling the activation of the I3 muscle).

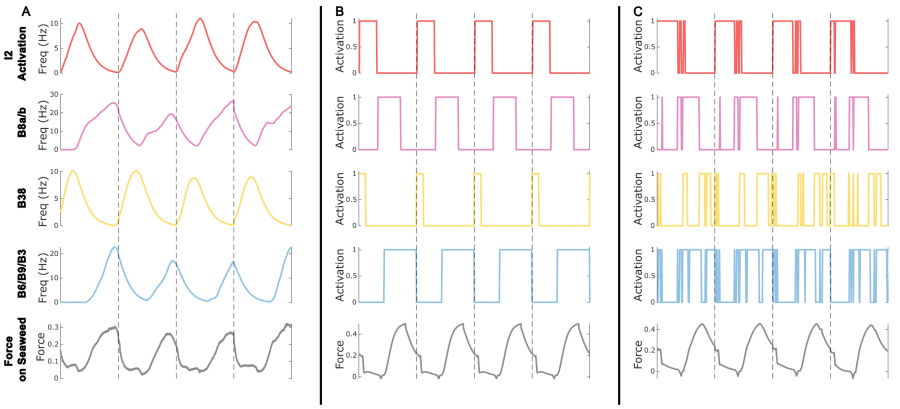

Fig. 1. Example swallowing performance datasets for (A) *Aplysia in vivo* experiments [9], (B) Expert model tuned behavior [13], and Trained model learned behavior [16]. The dashed lines separate consecutive swallowing cycles and were found using the segmentation and normalization methods. The separate swallows were individually used to make cross-correlation comparisons.

Segmentation and Normalization. Since we are interested in how different the models are to the animal data, in other words, how well the models correlate to animal data, we wanted to segment and normalize the animal data, Expert model data, and Trained model data to compare swallows with different phase durations to each other. For the segmentation, we split our time series data based on the onset of B31/32 neuron activation [9]. B31/32 innervates the I2 protractor muscle, rapidly decreasing the force on the seaweed after force peaks during the

retraction stage of the previous swallow [17]. The odontophore is pushed to the jaw opening, releasing the tension on the seaweed. In the reinforcement learning model, multiple B31/B32 activity episodes may occur within the same swallow. Therefore, to distinguish these bursts from the initial onset, a second segmentation criterion was introduced so that the segmentation should only happen if both an onset of B31/B32 activity and a steep negative slope in the force data were observed. Following segmentation, each swallow was time-normalized before subsequent comparisons. For normalization, time was divided by the total time elapsed in the segment, scaling all swallowing cycles to a normalized time of one, allowing for direct swallow comparisons. After applying the segmentation and normalization described (Fig. 1), the sample size for the animal swallow dataset is N = 24, the expert predicted swallow dataset is N = 1, and the trained predicted swallow dataset is N = 26. Note that the expert dataset is a single swallow because the neural model activation does not vary with every swallow cycle, making each swallow in a sequence of swallows exactly the same.

2.2 Statistical Analysis

Animal and Model Performance Correlation. To assess the similarity between two signals, in this case, animal swallowing behavior and model-predicted swallowing behavior, we conducted a cross-correlation analysis. This involved shifting the animal data (x) by various lag values (m) and computing the dot product with the Expert or Trained model data (y). The resulting cross-correlation coefficients (c) quantify the degree of correlation between the two signals across different temporal displacements, described by the equation:

$$c = \max(\widehat{R}_{xy}(m - P)), \quad \text{where } m = 1, 2, ..., 2P - 1 \tag{1}$$

where $\widehat{R}_{xy}(m)$ is defined as

$$\widehat{R}_{xy}(m) = \begin{cases} \sum_{n=0}^{P-m-1} x_{(n+m)} y_n^* & m \geq 0 \\ \widehat{R}_{xy}^*(m) & m < 0 \end{cases} \tag{2}$$

where P is the amount of data point in both signals and the asterisk denotes complex conjugation. We want the max $\widehat{R}_{xy}(m-P)$ because that is the moment when the lag shift between the two signals results in the highest correlation [18]. This value then has to be normalized to be between [−1,1] so it is normalized as

$$c_{coeff} = \frac{c}{\widehat{R}_{xx}(0)\widehat{R}_{yy}(0)} \tag{3}$$

where Rxx(0) and Ryy(0) refer to the cross-correlation of one signal compared to itself with zero lag. Cross-correlation coefficients were calculated between the swallows of a single animal and the swallows of all animals to see how much Intra- and Inter-animal variation is typical. Additionally, we calculated the cross-correlation coefficients for all animal swallows to all Expert and Trained model

swallows. Finally, we calculated the cross-correlation coefficients between Expert and Trained models to see how different their simulated swallows are. We will call these five different correlations "Intra-Animal" correlations, "Inter-Animal" correlations, "Animal-to-Expert" correlations, "Animal-to-Trained" correlations, and "Expert-to-Trained" correlations, respectively.

Bootstrapped Confidence Interval. Because our data sets include multiple swallows from each individual (both real and modeled) as well as dependency within our time series data between motor activity and grasping force, statistical comparisons of correlation coefficient distributions were performed using block bootstrapping and confidence intervals. Block bootstrapping is a bootstrapping method used for time series dependence [19] that creates blocks of data with strong dependence and assumes data outside the blocks hold insignificant dependencies. In the case of the swallowing data in this work, we consider a single swallow to be a block and assume that each swallow can be considered sufficiently independent of the next and previous swallow for the purposes of our pipeline. For all comparisons, we applied block bootstrapping and randomly resampled 10,000 times the median difference between the two groups being compared, thereby creating histograms of the bootstrapped differences between the medians of these distributions. Using a 95% confidence interval, the difference between each group is significant if a median difference of 0 does not fall within the confidence interval (i.e., we can reject the null hypothesis) and non-significant if 0 lies within the confidence interval (i.e., we fail to reject the null hypothesis).

We use this bootstrapped confidence interval in three different ways: (1) for animal-to-model correlations (Animal-to-Expert and Animal-to-Trained) relative to the Inter-animal group, (2) for model to RL model correlations (Expert-to-Trained) relative to Inter-animal, and (3) for animal-to-model correlation comparisons (Animal-to-Expert relative to Animal-to-Trained). We will describe these three cases in more detail. (1) If the difference between the animal-to-model and Inter-animal correlations are significant for any given property, we hypothesize that the model did not accurately capture animal swallowing behavior for demand-driven complexity. However, if there is no significant difference, the model can capture animal swallowing behavior, so no improvement is needed for the tested model property. (2) To determine if there is high variability between the Expert and Trained model performance, indicating discrepancies between the tuned neural model's and the RL controller's neural activity and/or swallowing force, we assess whether the difference between Expert-to-Trained and Inter-animal correlations is significant. And lastly, (3) when the confidence interval between Animal-to-Expert and Animal-to-Trained is found, if there is no significant difference, we hypothesize that the biomechanical model should be prioritized since the performance accuracy with different neural models did not change. If the confidence interval does show a significant difference, then there are two possibilities. The neural model should be prioritized if the Animal-to-Trained correlation is higher than the Animal-to-Expert since replacing the neu-

ral model with an RL controller increases performance accuracy. Inversely, the biomechanical model should be prioritized if the Animal-to-Expert correlation is higher than the Animal-to-Trained correlation since having the tuned neural model increases performance accuracy. Using these three results, we can confidently pinpoint exactly which components in the neuromechanical model need to be prioritized for demand-driven complexity improvement in future, more advanced iterations.

3 Results and Discussion

After conducting the analysis pipeline on the Expert model for the selected neural activations (i.e., I2 Muscle Activation, B8a/b, B38, and B6/B9/B3) and the resulting force on unbreakable seaweed, we found that force could confidently capture animal-like swallowing behavior, while neural activations could not (Fig. 2). The cross-correlation coefficient distributions were first found for comparisons of Intra- and Inter-animal swallows to serve as a benchmark for the typical variability seen *in vivo*. The cross-correlation coefficient distributions for Intra-animal and Inter-animal comparisons had medians of 0.97 (IQR 0.96–0.98) and of 0.96 (IQR 0.93–0.97), respectively (Fig. 2). As outlined in Sect. 2.1, comparisons were then performed with cross-correlation coefficient distributions of all animal swallows compared to all Expert model swallows (Animal-to-Expert), all animal swallows compared to all Trained model swallows (Animal-to-Trained), and all Expert model swallows compared to all trained model swallows (Expert-to-Trained). The corresponding median and interquartile range were 0.95 (IQR 0.94–0.96), 0.95 (IQR 0.93–0.97), and 0.95 (IQR 0.95–0.96), respectively (Fig. 2). After calculating the bootstrapped confidence intervals of force for both the Expert and Trained models, the null hypothesis failed to be rejected for a confidence level of 95%, ranging from [−0.0011, 0.0187] (Fig. 3A) and [−0.0014, 0.0121] (Fig. 3B), respectively. Given that the expert-tuned model used in this work to test our analysis pipeline was hand-tuned to mimic the force profile from *Aplysia* swallowing data, it is unsurprising that the force time series was not statistically significantly different from that of the animal data.

On the other hand, the analyzed neural activities of both Expert and Trained models failed to accurately capture animal swallowing behavior. The cross-correlation coefficient distribution median and interquartile ranges for all modeled motor neuron activity relative to animal swallows is detailed in Table 1. Compared to the previously discussed force correlations, the correlations found for motor neuron activity are much more variable, with cross-correlation coefficients below 0.9, indicating lower model accuracy. The bootstrapped confidence intervals comparing both the Animal-to-Expert and Animal-to-Trained distributions relative to the Inter-animal distribution rejected the null hypothesis for all neural activities (Fig. 4, Table 2). This indicates that the correlation coefficients varied significantly from the expected distributions across multiple individual animals, and therefore, neither model accurately captures animal-like swallowing

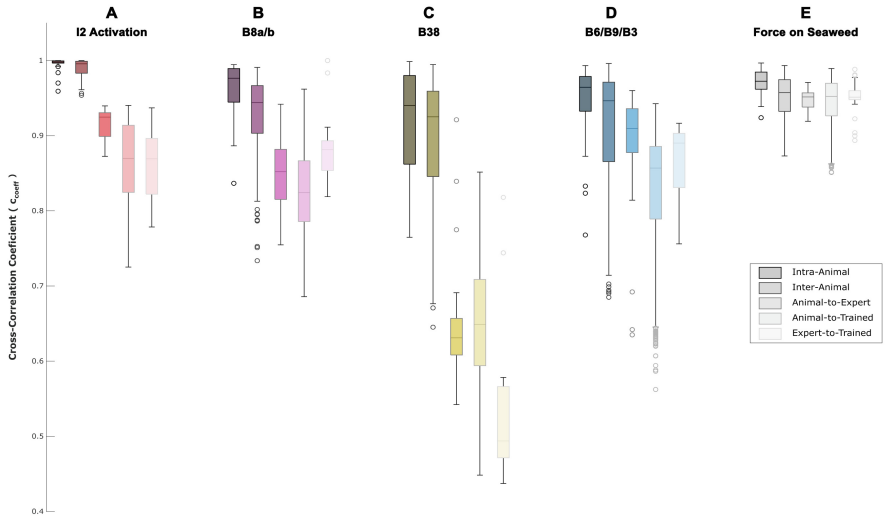

Fig. 2. Cross correlation distributions of animal and model comparisons. Boxes represent the 2nd and 3rd interquartile range with the center line at the median. Whiskers represent the 1st and 4th quartiles. Empty circles are outliers. Animal swallowing comparisons (both Intra-Animal and Inter-Animal) were performed for benchmarking, animal-to-model comparisons (Animal-to-Expert and Animal-to-Trained), and model-to-model comparisons (Expert-to-Trained) were performed as part of the analysis pipeline. Note that the sample size is not the same for the animal swallow dataset (N = 24), expert predicted swallow dataset (N = 1), and trained predicted swallow dataset (N = 26).

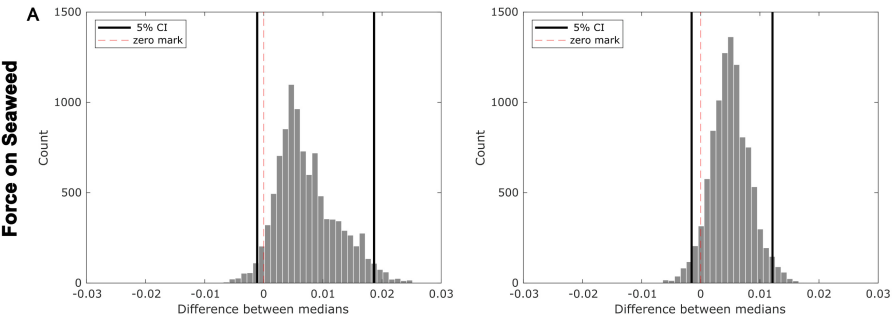

Fig. 3. Bootstrapping confidence intervals for force on seaweed for Animal-to-Expert (A) and Animal-to-Trained (B) relative to the Inter-animal group. Zero was found in both confidence intervals, represented by the red dashed line. Therefore, we fail to reject the null hypothesis. Thus, there is no significant difference between animal swallowing behavior and model-predicted animal behavior. (Color figure online)

Table 1. Median and interquartile range (IQR) of cross-correlation coefficient distributions for internal neural activation properties.

	I2 Activation	**B8a/b**	**B38**	**B3/B6/B9**
	Median	Median	Median	Median
	IQR	IQR	IQR	IQR
Intra-Animal	1.00	0.98	0.94	0.96
	1.00 - 1.00	0.94 - 0.99	0.86 - 0.98	0.93 - 0.98
Inter-Animal	0.99	0.94	0.92	0.95
	0.98 - 1.00	0.90 - 0.97	0.84 - 0.96	0.86 - 0.97
Animal-to-Expert	0.92	0.85	0.63	0.91
	0.90 - 0.93	0.81 - 0.88	0.61 - 0.66	0.88 - 0.93
Animal-to-Trained	0.87	0.82	0.65	0.86
	0.82 - 0.91	0.78 - 0.87	0.59 - 0.71	0.79 - 0.88
Expert-to-Trained	0.87	0.88	0.49	0.89
	0.82 - 0.90	0.85 - 0.89	0.47 - 0.57	0.83 - 0.90

activity for the neurons compared in this work. Alone, this information indicates that model improvements are needed, but we still need to determine specifically what aspects of the model should be targeted for demand-driven complexity improvements.

Table 2. Confidence interval (CI) for Expert and Trained model predictions for internal neural activation properties.

	I2 Activation	**B8a/b**	**B38**	**B6/9/3**
Expert CI	0.0664, 0.0841	0.0716, 0.1197	0.2705, 0.3102	0.0123, 0.0546
Trained CI	0.1207, 0.1295	0.1105, 0.1266	0.2573, 0.2929	0.0816, 0.1009

Comparing the Animal-to-Expert and Animal-to-Trained cross-correlation distribution to each other (Fig. 5), as well as comparing the Expert-to-Trained distribution to the Inter-animal distribution (Fig. 6), gives us insight into which portions of the model, either the neural model or biomechanical model, may need to be prioritized for improvement. For I2 muscle and motor neuron B6/B9/B3 activation, the confidence interval for the median difference between the Animal-to-Expert and Animal-to-Trained correlation distributions rejected the null hypothesis, the intervals being $[-0.0603, -0.0387]$ (Fig. 5A) and $[-0.0785, -0.0359]$ (Fig. 5D), respectively. Even though the Expert and Trained swallows are highly correlated to each other for these two properties (Table 1, Fig. 2A, 2D), there is a significant difference between how well they correlate with the animal swallowing dataset, with the Expert model having a higher correlation to animal swallowing than the Trained model (Fig. 5).

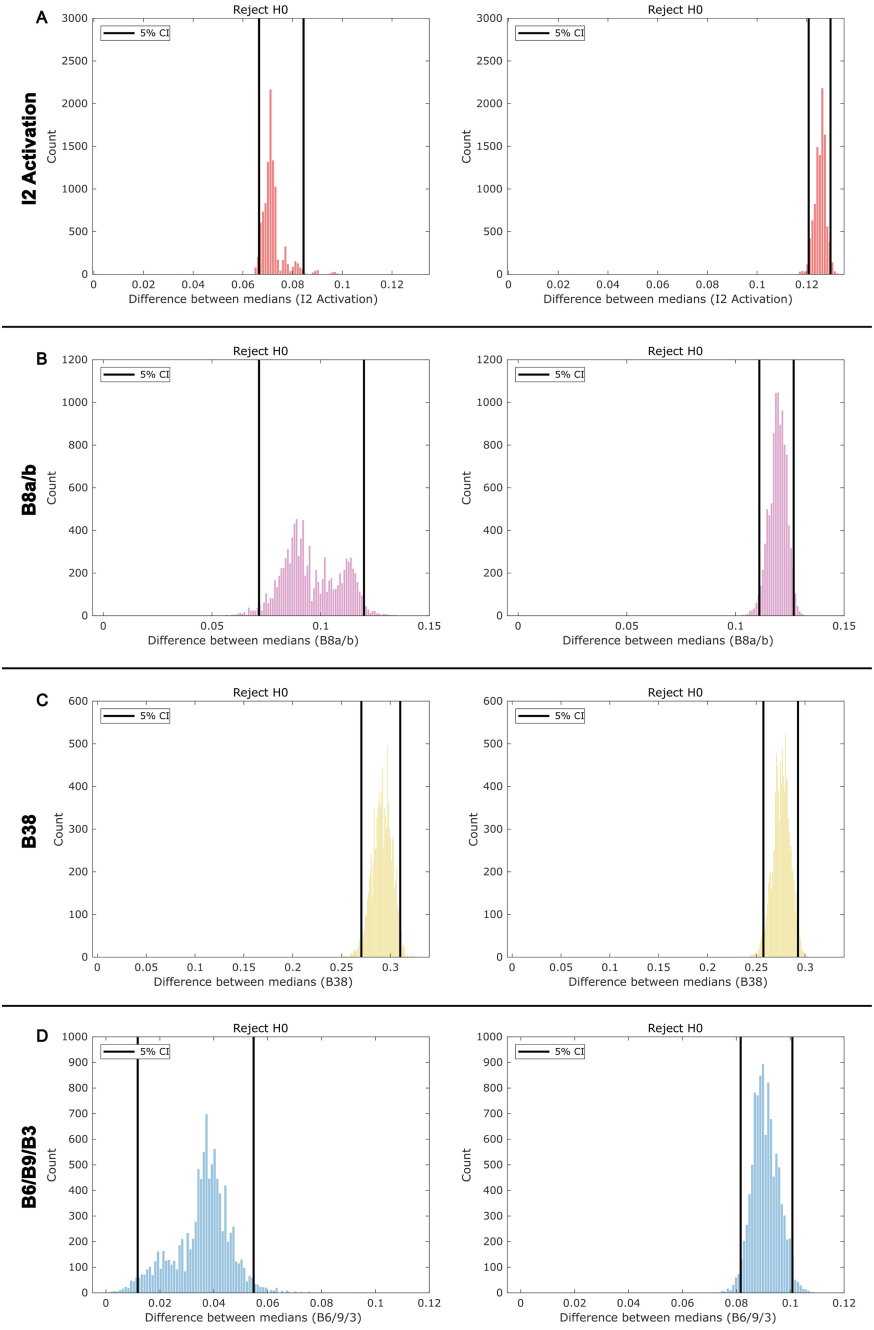

Fig. 4. Bootstrapping confidence intervals for motor neurons. A: I2 muscle activation, B: motor neuron B8a/b, C: motor neuron B38, and D: motor neurons B6/B9/B3 for Animal-to-Expert (left) and Animal-to-Trained (right) relative to the Inter-animal group. (A–D) We reject the null hypothesis for all activation properties since zero was not found inside any of the confidence intervals.

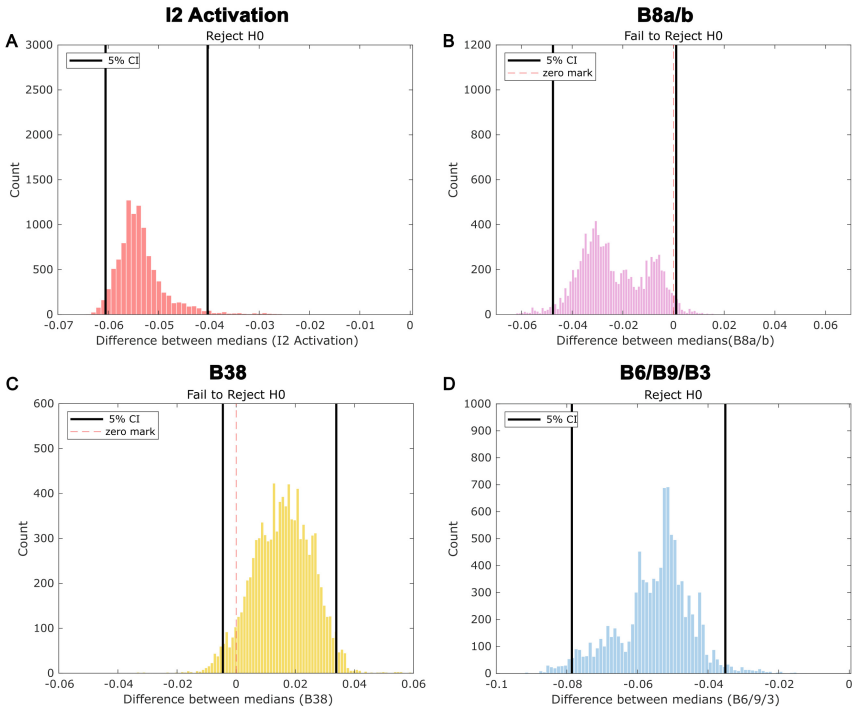

Fig. 5. Bootstrapping confidence intervals for motor neurons to determine which of the two components of the neuromechanical model, the neural or biomechanical model, needs to be prioritized for improvement. A: I2 muscle activation, B: motor neuron B8a/b, C: motor neuron B38, and D: motor neuron B6/B9/B3 for Animal-to-Expert relative to Animal-to-Trained. (A, D) We reject the null hypothesis since zero was not found inside the confidence interval. (B, C) Zero was found inside the confidence intervals, represented by the red dashed line. Therefore, we fail to reject the null hypothesis.

For motor neuron B8a/b, the confidence interval for the median difference between the Animal-to-Expert and Animal-to-Trained correlation distributions failed to reject the null hypothesis, with the confidence interval being $[-0.0482, 0.0012]$ (Fig. 5B). Therefore, there was no significant difference between how well the two models captured *in vivo* B8 activity in swallowing behavior. Based on this analysis, the RL policy did not learn I2, B3/B6/B9, or B8a/b activity that better correlated with animal data than the expert-tuned neuromechanical model. Therefore, future modeling efforts should focus on improving the biomechanical model associated with each of these motor neurons.

Whereas the trends for I2, B3/B6/B9, and B8a/b activity correlations were similar, all correlations of B38 stand out as being substantially lower than those of the other motor neurons. The confidence intervals comparison of the Animal-to-Expert and Animal-to-Trained cross-correlation distributions to that

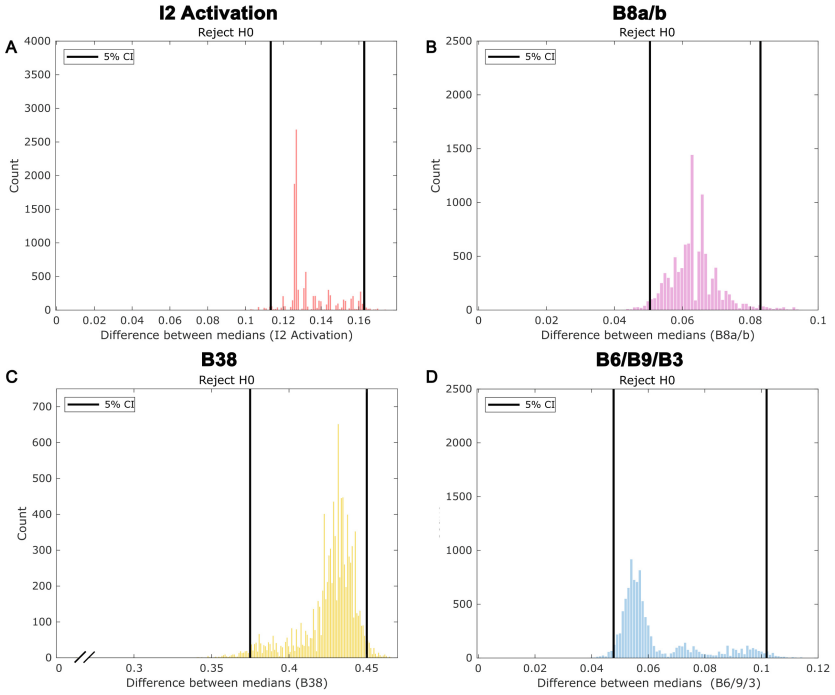

Fig. 6. Bootstrapping confidence intervals for motor neurons' performance correlation. A: I2 muscle activation, B: motor neuron B8a/b, C: motor neuron B38, and D: motor neurons B6/B9/B3 for Expert-to-Trained relative to the Inter-animal group. (A–D) We reject the null hypothesis for all activation properties since zero was not found inside any of the confidence intervals.

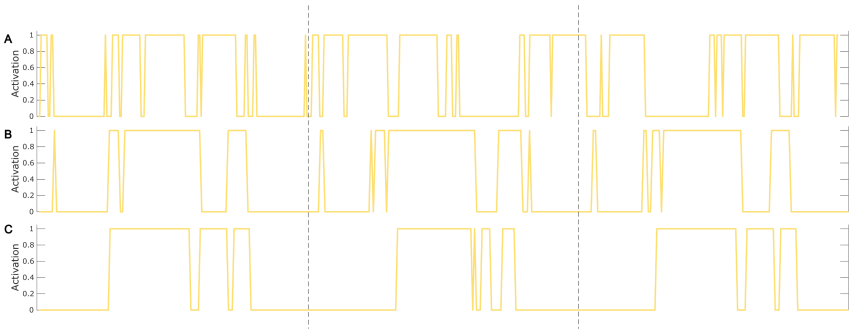

Fig. 7. Variability between individual trials of Trained model predicted swallow behavior for B38. The dashed lines separate consecutive swallowing cycles were found using the segmentation and normalization methods. A, B, and C are B38 boolean activation performances during three randomly selected trials, each swallow looking drastically different from consecutive swallows from the same trial and individual swallows from the other trials.

of Inter-animal variability of B38 were significant with [0.2705, 0.3102] and [0.2573, 0.2929], respectively (Fig. 4C). Furthermore, much wider variability was observed in the correlations for the trained B38 activity relative to the Inter-animal group (Table 1). Finally, comparing the Animal-to-Expert and Animal-to-Trained cross-correlation distribution to each other resulted in a confidence interval of [−0.0035, 0.0339] (Fig. 5C), indicating that not only did the models not correlate well with the animal data, they also did not correlate well with the animal data by a similar amount, with no significant difference in their performance accuracy, pointing to priority in improving the biomechanical model. *In vivo*, B38 activity in swallowing begins before I2 activation and then substantially overlaps I2 activation (Fig. 1A) [20]. In the Expert model, B38 activity is much shorter and occurs at the onset of I2 activation (Fig. 1B), whereas in the Trained model, B38 activity is highly variable, and much of the activity often occurs during the retraction phase activation of B6/B9/B3 and B8a/b (Figs. 7, 1C). Thus, the Expert and Trained models both exhibit B38 activity that is very different from the real animal behavior, and the two models are even more different from each other, with an Expert-to-Trained correlation median of 0.49 (IQR 0.47–0.57) and a confidence interval comparing Expert-to-Trained and Inter-animal variability showing a significant difference with [0.3749, 0.4501] (Table 1, Figs. 2C, 6C). Based on these outcomes, future model efforts should revisit both the biomechanics of pinch in the anterior region of I3 and the neural model implementation of B38. In the current biomechanical model, B6/B9/B3 also contribute to the jaw pinch [13,16], whereas the jaw pinch is likely more specifically mediated by B38 in the biomechanics of the real animal. This may make the B38 activation less important in the model implementations and, therefore, more prone to unrealistic variability. Additionally, jaw pinch during the retraction phase could hinder successful behavior by resisting inward seaweed movement; the presence of a large amount of B38 activation in the retraction phase of many Trained model responses suggests the model may not fully capture this.

4 Conclusions

This work presents a pilot study on a reinforcement-learning enabled analysis pipeline that classifies if an engineered neuromechanical model can confidently capture animal behavior, our test case being an existing model of *Aplysia californica* feeding. Using our analysis pipeline, we assessed this model to identify areas for future model refinement. We replaced the neural model with a reinforcement learning neural network and compared the correlation distributions of key motor neurons to those in animal swallowing data. By comparing the model's performance under its own neural control to that under RL-control, we showed that despite the model being able to capture swallowing force, the biomechanical model may not fully capture the complexities of *in vivo* neuromechanics. Furthermore, in comparing the correlations of the neural controller to the RL controller relative to Animal-to-Animal variability, we see that they are markedly

different and yet can both generate similar force outputs. This finding may indicate that the biomechanical model does not fully capture the mechanics of the *Aplysia* feeding apparatus, especially concerning motor neuron B38. Another limitation of the model being analyzed in this work is that the Boolean neural model is constrained to be either on or off, whereas real neurons can show much more graded spiking frequencies. The rapid alternation of on/off states in the Boolean model may greatly reduce the correlations, whereas using a more continuous model (e.g., an integrate-and-fire model [21], a synthetic nervous system model [22], or a multi-conductance Hodgkin-Huxley-like model [21]) might show less variation and thus yield better correlations. Finally, the analysis pipeline is limited by the small amount of experimental animal data available. A larger sample size would capture a wider range of possible animal behaviors and yield more precise results.

Future work should compare the results of this analysis pipeline when a more advanced biomechanical model is implemented to determine if this RL-enabled analysis approach has accurately identified areas for improvement. More experimental animal data can also be collected to have a larger sample of *Aplysia* swallowing behavior. Additionally, future tools should supplement this pipeline with an additional RL-enabled analysis wherein the neural model serves as the environment, and the biomechanical model is replaced with an RL policy, thereby further classifying areas for future demand-driven complexity model enhancement.

Acknowledgements. The authors thank Dr. Jeff Gill for sharing the animal data from his 2020 publication [9] for use in these comparisons. We thank the anonymous reviewers for their insightful edits and suggestions.

References

1. Karakasiliotis, K., et al.: From cineradiography to biorobots: an approach for designing robots to emulate and study animal locomotion. J. R. Soc. Interface **13**(119), 20151089 (2016). https://doi.org/10.1098/rsif.2015.1089
2. Ijspeert, A.J.: Biorobotics: using robots to emulate and investigate agile locomotion. Science **346**(6206), 196–203 (2014). https://doi.org/10.1126/science.1254486
3. Ostrowski, J., Burdick, J.W.: The geometric mechanics of undulatory robotic locomotion. Int. J. Robot. Res. **17**(7), 683–701 (1998). https://doi.org/10.1177/027836499801700701
4. Aguilar, J., et al.: A review on locomotion robophysics: the study of movement at the intersection of robotics, soft matter and dynamical systems. Rep. Prog. Phys. **79**(11), 110001 (2016). https://doi.org/10.1088/0034-4885/79/11/110001
5. Sun, X., et al.: A neuromechanical model for Drosophila larval crawling based on physical measurements. BMC Biol. **20**(1), 130 (2022). https://doi.org/10.1186/s12915-022-01336-w
6. Markin, S.N., Klishko, A.N., Shevtsova, N.A., Lemay, M.A., Prilutsky, B.I., Rybak, I.A.: Afferent control of locomotor CPG: insights from a simple neuromechanical model. Ann. N. Y. Acad. Sci. **1198**(1), 21–34 (2010). https://doi.org/10.1111/j.1749-6632.2010.05435.x

7. Sreenivasa, M., Valero-Cuevas, F.J., Tresch, M., Nakamura, Y., Schouten, A.C., Sartori, M.: Editorial: neuromechanics and control of physical behavior: from experimental and computational formulations to bio-inspired technologies. Front. Comput. Neurosci. **13**, (2019). https://doi.org/10.3389/fncom.2019.00013
8. Higuchi, K., Kazawa, T., Sakai, B., Namiki, S., Haupt, S.S., Kanzaki, R.: High performance, large-scale multi-compartment Hodgkin-Huxley simulation of Drosophila's whole-brain neural circuit model. bioRxiv (Cold Spring Harbor Laboratory) (2022). https://doi.org/10.1101/2022.11.01.512969
9. Gill, J.P., Chiel, H.J.: Rapid adaptation to changing mechanical load by ordered recruitment of identified motor neurons. eneuro, pp. ENEURO.0016-20.2020 (2020). https://doi.org/10.1523/eneuro.0016-20.2020
10. Lyttle, D., Gill, J., Shaw, K.M., Thomas, P., Chiel, H.J.: Robustness, flexibility, and sensitivity in a multifunctional motor control model, vol. 111, no. 1, pp. 25–47 (2016). https://doi.org/10.1007/s00422-016-0704-8
11. Shaw, K.M., et al.: The significance of dynamical architecture for adaptive responses to mechanical loads during rhythmic behavior. J. Comput. Neurosci. **38**(1), 25–51 (2014). https://doi.org/10.1007/s10827-014-0519-3
12. Novakovic, V.A., Sutton, G.P., Neustadter, D.M., Beer, R.D., Chiel, H.J.: Mechanical reconfiguration mediates swallowing and rejection in *Aplysia californica*. J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol. **192**(8), 857–870 (2006). https://doi.org/10.1007/s00359-006-0124-7
13. Webster-Wood, V.A., Gill, J.P., Thomas, P.J., Chiel, H.J.: Control for multifunctionality: bioinspired control based on feeding in *Aplysia californica*. Biol. Cybern. **114**(6), 557–588 (2020). https://doi.org/10.1007/s00422-020-00851-9
14. Li, Y., Sukhnandan, R., Gill, J.P., Chiel, H.J., Webster-Wood, V., Quinn, R.D.: A bioinspired synthetic nervous system controller for pick-and-place manipulation. IEEE Int. Conf. Robot. Autom. (ICRA) **2023**, 8047–8053 (2023). https://doi.org/10.1109/icra48891.2023.10161198
15. Dai, K., et al.: SLUGBOT, an *Aplysia*-inspired robotic grasper for studying control. In: Lecture Notes in Computer Science, pp. 182–194 (2022). https://doi.org/10.1007/978-3-031-20470-8_19
16. Sun, W., Xu, M., Gill, J., Thomas, P., Chiel, H.J., Webster-Wood, V.A.: GymSlug: Deep Reinforcement Learning Toward Bio-inspired Control Based on *Aplysia californica* Feeding, pp. 236–248 (2022). https://doi.org/10.1007/978-3-031-20470-8_24
17. Hurwitz, I., Kupfermann, I., Susswein, A.J.: Different roles of neurons B63 and B34 that are active during the protraction phase of buccal motor programs in *Aplysia californica*. J. Neurophysiol. **78**(3), 1305–1319 (1997). https://doi.org/10.1152/jn.1997.78.3.1305
18. "Cross-correlation - MATLAB xcorr. www.mathworks.com. https://www.mathworks.com/help/matlab/ref/xcorr.html
19. Kreiss, J.P., Paparoditis, E.: Bootstrap methods for dependent data: a review. J. Korean Stat. Soc. **40**(4), 357–378 (2011). https://doi.org/10.1016/j.jkss.2011.08.009
20. McManus, J.M., Lu, H., Cullins, M.J., Chiel, H.J.: Differential activation of an identified motor neuron and neuromodulation provide *Aplysia's* retractor muscle an additional function. J. Neurophysiol. **112**(4), 778–791 (2014). https://doi.org/10.1152/jn.00148.2014

21. Yamazaki, K., Vo-Ho, V.-K., Bulsara, D., Le, N.: Spiking neural networks and their applications: a review. Brain Sci. **12**(7), 863 (2022). https://doi.org/10.3390/brainsci12070863
22. Li, Y., Webster-Wood, V.A., Gill, J.P., Sutton, G.P., Chiel, H.J., Quinn, R.D.: A synthetic nervous system controls a biomechanical model of *Aplysia* feeding. In: Lecture Notes in Computer Science, pp. 354–365 (2022). https://doi.org/10.1007/978-3-031-20470-8_35