Synergizing Anatomy and Function: A Goal-driven Model of Frontoparietal Dexterous Object Manipulation

Tonio Weidler (t.weidler@maastrichtuniversity.nl)
Department of Cognitive Neuroscience, Faculty of Psychology and Neuroscience, Maastricht University, P.O. Box 616
Maastricht, 6200 MD, Limburg, The Netherlands

Rainer Goebel (r.goebel@maastrichtuniversity.nl)
Department of Cognitive Neuroscience, Faculty of Psychology and Neuroscience, Maastricht University, P.O. Box 616
Maastricht, 6200 MD, Limburg, The Netherlands

Mario Senden (mario.senden@maastrichtuniversity.nl)
Department of Cognitive Neuroscience, Faculty of Psychology and Neuroscience, Maastricht University, P.O. Box 616
Maastricht, 6200 MD, Limburg, The Netherlands

Abstract

Goal-driven deep learning produced significant advances in perception modelling. Models, however, often implement a single sensory domain and thus isolate a specific function. In this work, we go a step beyond and close the perception-action loop with a model of the frontoparietal network. The model implements biologically plausible macro-level structure by connecting cell count-fitted sensorimotor regions by pathways extracted from structural connectivity data. The model interfaces an anthropomorphic robotic hand and is trained to manipulate objects. We show that the biologically-inspired architecture significantly outperforms an architecture used in state-of-the-art robotics while converging substantially faster and relying only on raw sensory data. Moreover, preliminary in silico decoding analyses show promise in aligning with in vivo expectations.

Keywords: goal-driven modelling; sensorimotor control; dexterity; deep learning; reinforcement learning

Introduction

Understanding the neurocomputations that underlie human sensorimotor control is critical for elucidating the principles by which the human sensorimotor system functions. Of particular interest is human dexterity; the skilled and precise movements of hand and fingers that are unique to our species (van Leeuwen, Vanhoof, Kerkhof, Stevens, & Vereecke, 2018; Johnson-Frey, 2004). By examining how the brain processes sensory information in vast neural networks to coordinate movements during dexterous tasks, we can gain insights into the basic principles of motor control. The complexity and scale of these networks, however, presents a significant challenge for both data-driven and hand-engineered approaches. Building goal-driven models can serve as an alternative means of hypothesis testing and generation (Yamins & DiCarlo, 2016). Here, we present a bioinspired goal-driven model of the frontoparietal and pericentral networks (FPN) capable of manipulating objects with an anthropomorphic robotic hand.

Methods

The premise of goal-driven modelling is that under sufficient biological constraints, the neurocomputational strategies emerging in silico will be reminiscent of their biological counterparts. We apply three structural constraints when building our model: *region selection*, *capacity scaling*, and *pathway construction*, which we describe in the following paragraphs.

Region Selection Frontoparietal networks (FPN) are widely regarded as involved in the sensory-guided control of the human hand (Filimon, 2010; Binkofski et al., 1999; Ptak, Schnider, & Fellrath, 2017; Gallivan, Johnsrude, & Flanagan, 2015; Marek & Dosenbach, 2018). Comprised primarily of regions in posterior parietal cortex (PPC) and prefrontal cortex (PFC), the FPN bridges between sensory and motor cortices. Following the general taxonomy of Uddin, Yeo, and Spreng (2019), we model the lateral PFC and midcingulate cortex (MCC) in PFC, and the inferior and superior parietal lobules (IPL, SPL), and intraparietal sulcus (IPS) in PPC. Moreover, we model premotor cortex (PMC) and primary motor cortex (M1), as well as somatosensory cortex (SSC). These comprise the pericentral network (Uddin et al., 2019) and, in our model, wrap the FPN with inputs and outputs. Additionally, the object's pose is integrated as a surrogate latent visual representation (e.g. middle temporal complex MT/V5).

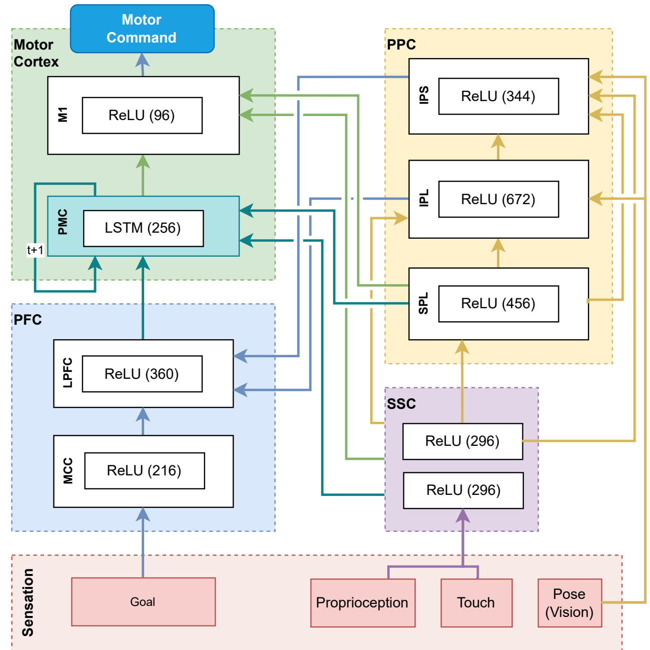

Figure 1: Our deep neural network architecture of the FPN. Activity vectors projected (indicated by arrows) to some region are concatenated for joined transformation.

For every region, we select the corresponding set of cytoarchitectonic regions from the Julich-Brain atlas (Amunts, Mohlberg, Bludau, & Zilles, 2020). In the following steps, a region is then represented by an aggregate of region-wise data.

Capacity Scaling Regions can vastly differ in their spatial extent and cellular density and, thus, in the number of neurons they comprise. Consequently, some regions impose informational bottlenecks on the network structure. By strategically placing bottlenecks along the model's pathways, the network's potential neurocomputational strategies can be significantly pruned, making them a crucial constraint.

A model region i's cell count c_i was extrapolated from cell densities reported in von Economo and Triarhou (2009), and region volumes provided by the Julich-Brain atlas (Amunts et al., 2020). Layer-wise densities were aggregated into region-wise densities using a weighted average that accounts for layer thickness. A mapping between Economo and Julich regions was established either based on direct matches, a Brodmann region map (Triarhou, 2007), or by manual assignment.

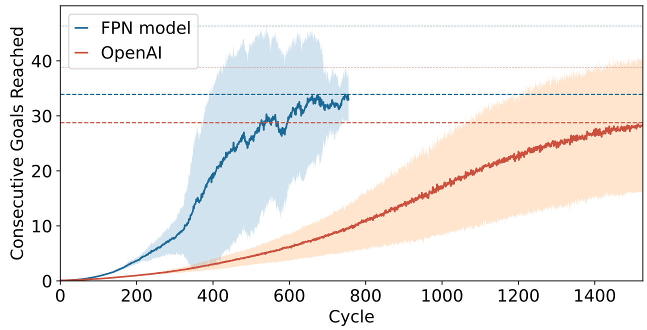

Figure 2: Training of a baseline architecture (OpenAI et al., 2020) compared to our own model. Both curves are averaged over 5 independent agents, and shaded regions indicate the 95% t-confidence interval about the mean. Horizontal lines mark the peak performance of the group mean (dotted) and of the best agent in each group (dashed).

Finally, a reference target in the PMC was set to $n_{pmc} = 256$, and all other region's unit counts n_i are scaled to $n_i = c_i \frac{n_{pmc}}{c_{pmc}}$. The results can be found in Figure 1.

Pathway Construction Together with each region's unit count, their connecting pathways will define how information can be processed to map from sensation to movement. Our model's pathways are biologically inspired by structural connectivity data (Domhof, Jung, Eickhoff, & Popovych, 2021). A normed connectivity matrix is thresholded at 0.2 to construct an undirected graph. The graph is then directed from sensation to motor, avoiding cycles and leafs other than M1.

Training The model is trained on the in-hand object manipulation task introduced in OpenAI et al. (2020), using AngoraPy (Weidler & Senden, 2023), a library for neuroscientific goal-driven sensorimotor modelling using deep reinforcement learning. We optimize a non-sparse reward for rotational progress toward a target orientation and augment it by a biologically inspired loss minimizing the squared applied force (Pedotti, Krishnan, & Stark, 1978).

Results

Figure 2 compares the learning of our FPN model and a baseline model implementing OpenAI et al. (2020)'s architecture applied to the same raw sensory states. The FPN model converges substantially faster than its non-biological counterpart, and to a higher performance. Until convergence, the agent observed approximately 1.5 years of experience, about half as many as the baseline. Table 1 shows the performance of the best run of each architecture. The FPN model significantly outperforms the baseline architecture by a 12.36% increase in achieved goals and can chain 48 reorientations.

Figure 3 shows preliminary in silico decoding results. Of all regions, only M1 and particularly PMC encode information about the reward, in line with in vivo findings (Ramkumar, Dekleva, Cooler, Miller, & Kording, 2016; Marsh, Tarigoppula, Chen, & Francis, 2015). Moreover, the PMC codes for information indicative of the number of goals already achieved within the current episode. It thus appears to track the valence of its history. PMC and M1 representations code for future object poses, but the M1's predictiveness deteriorates faster with temporal delay than the former's. This might indicate that PMC maintains memory *and* plans ahead. This fits neuroimaging results showing M1 and PMC to code for immediate action and entire sequences respectively (Yokoi, Arbuckle, & Diedrichsen, 2018).

Model	CGR (Sampling)	CGR (Mode)
FPN Architecture	46.33	48.06 ± 0.8
OpenAI Architecture	38.73	42.13 ± 1.35
OpenAI et al. (2020)	N/A	43.4

Table 1: Consecutive goals reached (CGR) by the best agents selected from the FPN and OpenAI et al. (2020) architecture. Performance is given both for the stochastic agent sampling actions during training and for an evaluation after training, where the agent selects the mode of the action distribution. The bottom row shows the performance reported in OpenAI et al. (2020), trained with several environment randomizations and preprocessed states.

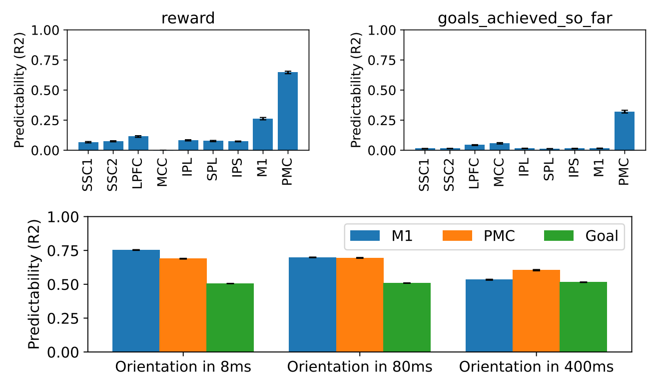

Figure 3: In silico decoding of three pieces of information from activity recorded in the model. In the bottom plot, predictions from the goal serve as a baseline.

Conclusion

We present a goal-driven FPN model with macro-scale plausibility. Our bioinspired architecture learns in-hand object manipulation and improves upon previous work in performance and sample efficiency. We thus contribute the first model of its kind implementing both biological plausibility and functional validity, serving a basis for research analysing its computations or building upon it. A preliminary analysis of the trained network's internal representations already shows promise with respect to the validity of emerging internal representations.

Acknowledgments

This study was funded by the European Union's Horizon 2020 Framework under the Specific Grant Agreement No. 945539 (HBP SGA3). We further acknowledge the use of Fenix Infrastructure resources, which are partially funded from the European Union's Horizon 2020 research and innovation programme through the ICEI project under the grant agreement No. 800858.

References

Amunts, K., Mohlberg, H., Bludau, S., & Zilles, K. (2020). Julich-Brain: A 3D probabilistic atlas of the human brain's cytoarchitecture. *Science*, *369*(6506), 988–992. doi: 10.1126/science.abb4588

Binkofski, F., Buccino, G., Posse, S., Seitz, R. J., Rizzolatti, G., & Freund, H. J. (1999). A fronto-parietal circuit for object manipulation in man: Evidence from an fMRI-study. *European Journal of Neuroscience*, *11*(9), 3276–3286. doi: 10.1046/j.1460-9568.1999.00753.x

Domhof, J. W. M., Jung, K., Eickhoff, S. B., & Popovych, O. V. (2021). *Parcellation-based structural and resting-state functional brain connectomes of a healthy cohort.* EBRAINS. doi: 10.25493/81EV-ZVT

Filimon, F. (2010). Human Cortical Control of Hand Movements: Parietofrontal Networks for Reaching, Grasping, and Pointing. *The Neuroscientist*, *16*(4), 388–407. doi: 10.1177/1073858410375468

Gallivan, J. P., Johnsrude, I. S., & Flanagan, J. R. (2015). Planning Ahead: Object-Directed Sequential Actions Decoded from Human Frontoparietal and Occipitotemporal Networks.
doi: 10.1093/cercor/bhu302

Johnson-Frey, S. H. (2004). The neural bases of complex tool use in humans. *Trends in Cognitive Sciences*, *8*(2), 71–78. doi: 10.1016/j.tics.2003.12.002

Marek, S., & Dosenbach, N. U. F. (2018). The frontoparietal network: Function, electrophysiology, and importance of individual precision mapping. *Dialogues in Clinical Neuroscience*, *20*(2), 133–140.

Marsh, B. T., Tarigoppula, V. S. A., Chen, C., & Francis, J. T. (2015). Toward an Autonomous Brain Machine Interface: Integrating Sensorimotor Reward Modulation and Reinforcement Learning. *Journal of Neuroscience*, *35*(19), 7374–7387. doi: 10.1523/JNEUROSCI.1802-14.2015

OpenAI, Andrychowicz, M., Baker, B., Chociej, M., Józefowicz, R., McGrew, B., ... Zaremba, W. (2020). Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, *39*(1), 3–20. doi: 10.1177/0278364919887447

Pedotti, A., Krishnan, V. V., & Stark, L. (1978). Optimization of muscle-force sequencing in human locomotion. *Mathematical Biosciences*, *38*(1), 57–76. doi: 10.1016/0025-5564(78)90018-4

Ptak, R., Schnider, A., & Fellrath, J. (2017). The Dorsal Frontoparietal Network: A Core System for Emulated Action. *Trends in Cognitive Sciences*, *21*(8), 589–599. doi: 10.1016/j.tics.2017.05.002

Ramkumar, P., Dekleva, B., Cooler, S., Miller, L., & Kording, K. (2016). Premotor and Motor Cortices Encode Reward. *PLOS ONE*, *11*(8), e0160851. doi: 10.1371/journal.pone.0160851

Triarhou, L. C. (2007). A Proposed Number System for the 107 Cortical Areas of Economo and Koskinas, and Brodmann Area Correlations. *Stereotactic and Functional Neurosurgery*, *85*(5), 204–215. doi: 10.1159/000103259

Uddin, L. Q., Yeo, B. T., & Spreng, R. N. (2019). Towards a universal taxonomy of macro-scale functional human brain networks. *Brain topography*, *32*(6), 926–942. doi: 10.1007/s10548-019-00744-6

van Leeuwen, T., Vanhoof, M. J. M., Kerkhof, F. D., Stevens, J. M. G., & Vereecke, E. E. (2018). Insights into the musculature of the bonobo hand. *Journal of Anatomy*, *233*(3), 328–340. doi: 10.1111/joa.12841

von Economo, C., & Triarhou, L. C. (2009). *Cellular Structure of the Human Cerebral Cortex.* Karger.

Weidler, T., & Senden, M. (2023). *AngoraPy - Anthropomorphic Goal-Oriented Robotic Control for Neuroscientific Modeling.* Zenodo. doi: 10.5281/zenodo.6636482

Yamins, D. L. K., & DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, *19*(3), 356–365. doi: 10.1038/nn.4244

Yokoi, A., Arbuckle, S. A., & Diedrichsen, J. (2018). The role of human primary motor cortex in the production of skilled finger sequences. *Journal of Neuroscience*, *38*(6), 1430–1442. doi: 10.1523/JNEUROSCI.2798-17.2017