UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Executive Function Variability in Biologically Plausible and Optimized Neural Agents

Permalink

https://escholarship.org/uc/item/8xt7g9vv

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors

Ezra Tsur, Elishai
Cohen Duwek, Hadar
Clark, Yisrael

Publication Date

2026

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Executive Function Variability in Biologically Plausible and Optimized Neural Agents

Yisrael Clark (yisrael.clark@gmail.com)
Hadar Cohen-Duwek (hadarco@openu.ac.il)
Elishai Ezra Tsur (elishai@nbel-lab.com)

Neuro-Biomorphic Engineering Lab (NBEL), Department of Mathematics and Computer Science
The Open University of Israel, Ra'anana, Israel

Abstract

Executive functions support flexible, goal-directed behavior in dynamic environments. Existing models typically attribute the substantial variability in human performance on such tasks to abstract, resource-based accounts of attention and working memory, leaving the nature of the underlying resource underspecified. Here, we investigate executive-function variability using a modified, continuous version of the Pong game that imposes sustained demands on multi-object tracking, control, and action selection. To account for the broad distribution of human performance levels, we compared a performance-optimized deep reinforcement learning agent with a biologically plausible spiking neural network that incorporates explicit neural resource constraints. Whereas the unconstrained artificial agent achieved superhuman performance, the spiking model reproduced the full spectrum of human scores. Systematic variation in neural resources, perceptual noise, and processing delays produced performance regimes corresponding to low, median, and peak human behavior, including saturation at object counts consistent with known working-memory limits.

Keywords: Neuromorphic Computing, Executive Function, Neural Engineering Framework, Computational Neuroscience, Reinforcement Learning, Spiking Neural Network, Multiple Object Tracking.

Introduction

Executive functions (EF) are higher-order cognitive processes underlying flexible, goal-directed behavior, including working memory updating, sustained attention, task switching, and inhibitory control. (Diamond, 2013). EF are most evident in continuous, dynamic environments that require ongoing monitoring and action selection. A central goal of computational models of EF is to develop mechanistic models that explain how neural circuits implement these functions and why human performance is limited and variable (Kriegeskorte & Douglas, 2018).

A key component of EF is the ability to maintain and manipulate multiple items in working memory. Studies of multiple object tracking (MOT) show that humans can typically track three to five objects simultaneously (Pylyshyn & Storm, 1988), reflecting limits in attentional selection. A similar capacity limit has been observed in visual working memory (VWM), which supports the retention of visual information after stimulus offset (Cowan, 2010). Although these limits are numerically similar, evidence suggests that attentional tracking and VWM rely on partially distinct mechanisms (Fougnie & Marois, 2006).

Early accounts explained these capacity limits in terms of a fixed number of discrete storage (Cowan, 2001; Zhang & Luck, 2008). More recent work, however, characterizes capacity as a limited but divisible resource that can be flexibly allocated across items (Bays and Husain, 2008; Ma et al., 2014). Performance declines gradually as the number of tracked objects increases, because fewer resources are available for representing each item (Alvarez & Franconeri, 2007). While this framework captures key behavioral phenomena, the neural substrate of the proposed "resource" remains largely underspecified.

Neuromorphic modeling offers a principled approach for grounding such resource constraints in biologically realistic neural mechanisms. Spiking neural networks (SNNs) operate through discrete spike events and temporal dynamics, capturing essential properties of biological neural computation while imposing concrete constraints on representation and processing (Maass, 1997; Tsur, 2021). These constraints provide a natural basis for linking neural architecture to observed limits and variability in cognitive performance.

From a computational perspective, reinforcement learning (RL) provides powerful methods for learning goal-directed behavior through interaction with complex environments (Sutton & Barto, 2018). Deep RL algorithms, such as Soft Actor-Critic, can achieve remarkably high performance on demanding control tasks (Haarnoja et al., 2018). However, these systems typically rely on non-biological architectures and unconstrained optimization, limiting their usefulness as explanatory models of human cognition.

Video games, such as the Atari 2600 games, provide a unique bridge between these domains. They serve both as standard benchmarks for RL research (Bellemare et al., 2013; Fan, 2023) and as ecologically valid tools for probing executive functions in dynamic, visually complex environments (Takahashi et al., 2024). Neuroimaging studies using Atari games show that human brain activity encodes abstract state representations analogous to those learned by deep RL agents (Cross et al., 2021). The game Pong has been especially influential, spanning applications from in vitro learning systems to human cognitive experiments (Blakowski et al., 2025; Kagan et al., 2022).

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society* ©2026 the author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

In this study, we develop a neuromorphic model of executive function using a modified Pong task that places continuous demands on multi-object tracking and action selection. Human performance on this task exhibits substantial interindividual variability. We first train a performance-optimized artificial neural network using state-of-the-art reinforcement learning methods, and then convert it into a fully spiking, biologically plausible network embedded within a larger SNN architecture. The model incorporates a capacity-limited object tracking mechanism, analogous to attentional and working memory systems, together with an action selection circuit inspired by cortico–basal ganglia–thalamocortical loops (Alexander et al., 1986). We show that, unlike the unconstrained ANN, the neuromorphic model reproduces the full range of human performance, with systematic variation linked to object tracking capacity, consistent with established limits in working memory. These results demonstrate how biologically grounded neural constraints can provide a mechanistic account of variability in executive function.

Methods and Materials

Human Behavioral Experiment

We developed a modified version of the Atari game Pong using the Unity game engine to assess executive function. The task was part of a broader set of short video game-based assessments, though only the Pong task was examined here. Participants were required to keep multiple balls in play by moving a horizontal paddle, engaging eye-hand coordination, multi-object tracking (MOT), inhibitory control, and flexible shifting of attention as task demands increased.

Participants controlled the paddle using left and right arrow keys, with paddle speed held constant throughout the task. Balls followed physics-based trajectories governed by gravity and collisions with the paddle, the environment boundaries, and other balls. Multiple balls could be present concurrently and interacted dynamically, increasing perceptual and control demands.

Task difficulty increased adaptively based on performance. One additional ball was introduced after every five successful paddle–ball contacts. When a ball was missed, it was immediately respawned, while the internal counter of succesful hits was reduced by half a point, preventing abrupt difficulty spikes during periods of low performance.

Each session lasted 2.5 minutes and included an initial 5-second delay before ball appearance. Participants rested briefly before gameplay onset. Performance was quantified as the cumulative number of successful paddle-ball contacts over the session.

Participants: Twenty-six participants (15 male, 11 female) aged 13 to 60 years (mean age of 34) were recruited for this study. All participants provided informed consent, and the study was approved by the institutional ethics committee. Each participant completed a single trial of the modified Pong game on a personal computer. All participants reported being healthy, with no history of diagnosed cognitive impairments.

Data Analysis. Behavioral data were analyzed using Python. The primary dependent variable was the score for each participant. Descriptive statistics, including the mean, standard deviation (SD), median, minimum, and maximum scores, were calculated to quantify the central tendency and variability of human performance.

Deep Reinforcement Model

To establish a benchmark of performance achievable through pure optimization, a standard ANN agent was developed and trained using deep reinforcement learning.

Network Architecture: The game environment interfaced with reinforcement learning algorithms via Unity's ML-Agents toolkit (Juliani et al., 2020). We used The Soft Actor-Critic (SAC) algorithm, which incorporates an actor model and a critic model: The actor selects actions based on a policy, and the critic evaluates those actions by estimating their value (expected future reward). The critic's feedback guides the actor to adjust its policy toward actions that yield higher long-term rewards, combining the strengths of policy-based and value-based learning (Konda & Tsitsiklis, 1999). The details of the model are as follows:

Input Representation: The observation space included the positions and velocities of up to 9 balls simultaneously present in the game, and the paddle location. This resulted in a high-dimensional observation vector capturing the complete state of the game at each timestep. When fewer balls were present, remaining slots were zero-padded, distinguishing inactive balls via invalid velocity.

Network Structure: The actor network architecture consisted of a single dense layer projecting the observation vector to 256 hidden units with softLIFRate activation functions. Leaky Integrate-and-Fire (LIF) neurons are widely used in neuromorphic and computational neuroscience models because they provide a simple yet biologically grounded approximation of real neuronal membrane dynamics, balancing computational efficiency with biological plausibility (Tsur, 2021). SoftLIFRate neurons are rate-based approximations of LIF neurons with smoothing around the firing threshold, designed to facilitate later conversion to spiking neural networks while maintaining compatibility with gradient-based learning (Hunsberger & Eliasmith, 2015). The softLIFRate parameters were set to $\tau_{RC} = 0.02s$ (membrane time constant) and $\tau_{ref} = 0.002s$ (refractory period), consistent with standard NEF implementations. The activations were scaled by 0.002 during training to produce a dynamic range similar to ReLU activations, the activiation orignally used by the SAC algorithm, thereby ensuring stable training with the SAC algorithm. The critic network employed three hidden layers of 256 neurons each to estimate state-action value functions. Since only the trained actor model was needed for the neuromorphic model, the critic model was trained using ReLU activations, as in the original paper (Haarnoja et al., 2018), and not softLIFRate activations.

Output: From the 256 hidden units of the actor network, two separate output layers were computed following the SAC

architecture: a mean (μ) layer and a standard deviation (σ) layer, each projecting to a single output neuron with tanh activation. These outputs parameterized a Gaussian policy over continuous actions (paddle movement).

Training: The ANN agent was trained within the game environment using the SAC algorithm. The agent received a sparse reward signal to guide its learning. A reward of +1 was given for each successful ball interception. No reward was given at any other time step. Key hyperparameters, such as the learning rate, discount factor $\boldsymbol{\gamma}$, and algorithm-specific parameters were selected based on common practices in the RL literature to ensure stable and effective learning. Given that each game session lasted 150 seconds and the environment updated every 0.1 seconds, each complete game comprised 1,500 timesteps. The agent was trained for 5 million time-steps, the equivalent of 3,333 games.

Initial training exploited an inhuman strategy: bouncing the balls off the very tip of the paddle and achieving artificially inflated scores. To prevent this, gaussian noise was added to the observations provided by the game, as found in humans (Michel & Geisler, 2011). This biologically plausible constraint induced more human-like playing in the model, with more reasonable scores, as detailed below.

Spiking Neural Network (SNN) Model

The primary computational model of this study is a biologically plausible SNN designed to explain, rather than merely optimize, performance. Its construction was guided by principles of neural computation and known brain architectures.

The Neural Engineering Framework (NEF) enables the construction of large-scale functional brain models using spiking neurons, translating computational objectives into biologically constrained neural architectures (Eliasmith & Anderson, 2003; Tsur, 2021). This framework has proven effective for modeling diverse cognitive functions, including visual processing (Tsur & Rivlin-Etzion, 2020), perception (Cohen-Duwek et al., 2022), and decision-making (Eliasmith et al., 2012), while maintaining biological plausibility. In this work, NEF was used, through its realization in the Nengo library (Bekolay et al. 2014), to encode task-relevant variables in neural populations, compute nonlinear transformations between populations, and implement simple dynamics through recurrent connections with biologically plausible synaptic time constants.

Network design. The SNN's architecture was designed to reflect a plausible cognitive and neural hierarchy, organized as a simplified cortico-basal ganglia-thalamocortical loop, a circuit known to be critical for decision-making and executive control (Redgrave et al., 1999). The model was organized into three functionally distinct but interacting subsystems, as seen in Figure 1, corresponding to key executive components: (1) a perceptual and working memory module that maintains and updates visual information from the Pong environment, (2) a cognitive processing module that integrates dynamic state information using a spiking neural implementation of a trained ANN, and (3) an action selection module modeled after the basal ganglia-thalamocortical circuitry responsible for response selection and inhibition (Stewart et al., 2010). Information flows from perception to the processing system (representing cortical processing), which then provides the relevant state information to the action selection system (representing the basal ganglia), whose output drives the motor response. All synaptic connections in the model employed time constants within the biologically plausible range reported for cortical neurons, typically 5–100ms (Destexhe et al., 2003).

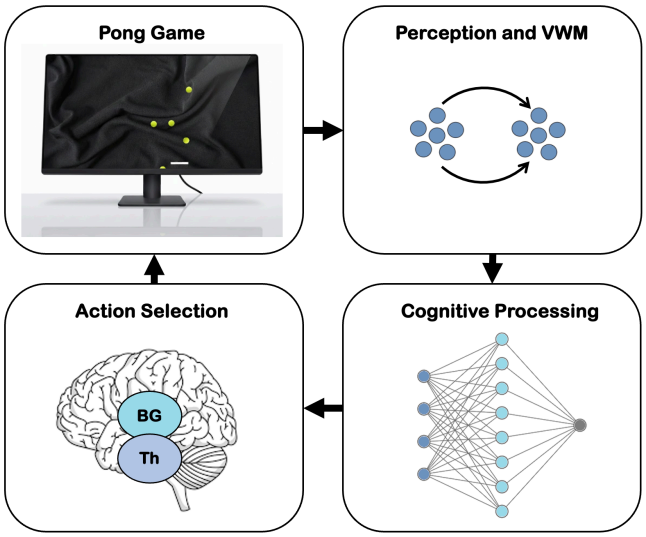

Figure 1: Neuromorphic model architecture. The Unity game environment provides observations. The Perception and VWM module use ensembles to represent and calculate the ball locations and velocities. The Cognitive Processing Module receives these representations and calculates action values using the converted ANN. The Action Selection Module chooses the action with the highest utility using a Basal Ganglia-Thalamus circuit and feeds the chosen action back into the Unity game environment.

Perception and Visual Working Memory. The perception module comprised two neural ensembles. One ensemble encoded ball positions, corresponding to attentional object tracking, while a second ensemble encoded ball velocities, supporting the storage and manipulation of object information for trajectory estimation. This separation reflects evidence that attention and visual working memory (VWM) are partially dissociable functions and can rely on distinct neural substrates (Fougnie & Marois, 2006; Roussy et al., 2021).

Ball positions were extracted from the game's observation vector via a linear transformation implemented in the connection weights. Although the game environment provides explicit velocity information, velocity was computed neuromorphically to enhance biological plausibility. Following Potjans et al. (2009), the velocity-encoding ensemble received position input through two parallel synaptic pathways: one with a short synaptic delay (τ

= 0.005 s) representing the current position, and one with an additional 0.1 s delay representing the previous position, matching the 100ms game update interval (Figure 2). Velocity was computed as the difference between these signals, effectively implementing a short-term memory mechanism for motion estimation. This computation can be interpreted as a component of VWM involved in motion processing (Kang et al., 2011).

The number of neurons and the dimensionality of each ensemble were systematically varied. Neuron count determined representational precision, while dimensionality constrained the number of objects that could be simultaneously encoded and tracked. These parameters were manipulated independently in both the attentional and VWM-related ensembles to examine how multi-object tracking limitations affected overall task performance.

Cognitive Processing Module. The actor component of the trained ANN was converted into a spiking neural network by constructing an architecturally matched model in Nengo. The 256 softLIFRate units in the trained actor were mapped one-to-one onto 256 LIF spiking neurons. Because softLIFRate activations were scaled by 0.002 during training, the output of the corresponding LIF ensemble was scaled by the same factor to preserve the learned representations. LIF neuron parameters were set to $\tau_{RC} = 0.02\,s$ and $\tau_{ref} = 0.002\,s$, matching the training configuration.

Unlike the rest of the model, which follows NEF principles of population encoding and decoding, the converted ANN neurons operate via rate encoding. Each spiking neuron corresponds directly to a single softLIFRate unit, with information represented by firing rate rather than by distributed population activity. Both population-based and rate-based encoding schemes are supported by prior theoretical and empirical work (Abbott & Dayan, 2001).

The converted network retained the dense hidden layer and the mean (μ) output layer, while the standard deviation (σ) layer was removed, as it is required only during training for stochastic policy sampling. The $tanh$ output nonlinearity was also removed, since the downstream basal ganglia action selection circuit constrained actions to discrete left, right, or no-movement outputs (Figure 2).

Action Encoding and Basal Ganglia Selection. Action selection was implemented using a neuromorphic model of the basal ganglia–thalamus loop, the brain's primary system for resolving competition between potential actions (Stewart et al., 2010). The continuous output of the converted ANN was transformed into discrete actions through this biologically inspired circuit, which provides a functional abstraction of basal ganglia pathways, including striatal input and direct and indirect (Go/NoGo) action selection mechanisms. The model was implemented using a standard Nengo basal ganglia module.

Three action channels were defined, corresponding to left movement, right movement, and no movement. Continuous ANN output was mapped to these channels using population encoding. The left and right movement ensembles employed directionally tuned neurons, with positive values preferentially activating the right-movement population and negative values activating the left-movement population. A constant threshold input was provided to the no-movement ensemble, biasing selection toward inaction when the ANN output was near zero. This mapping translated a continuous control signal into discrete action utilities, consistent with population-based motor encoding (Georgopoulos et al., 1986; Oryshchuk et al., 2024).

The resulting action utilities were provided as input to the basal ganglia network, configured for three-way action selection using default Nengo parameters. The thalamic output selected the highest-utility action while suppressing alternatives. The selected action was then communicated to the game environment via a Nengo interface to ML-Agents, updating the environment at 100 ms intervals (Figure 2).

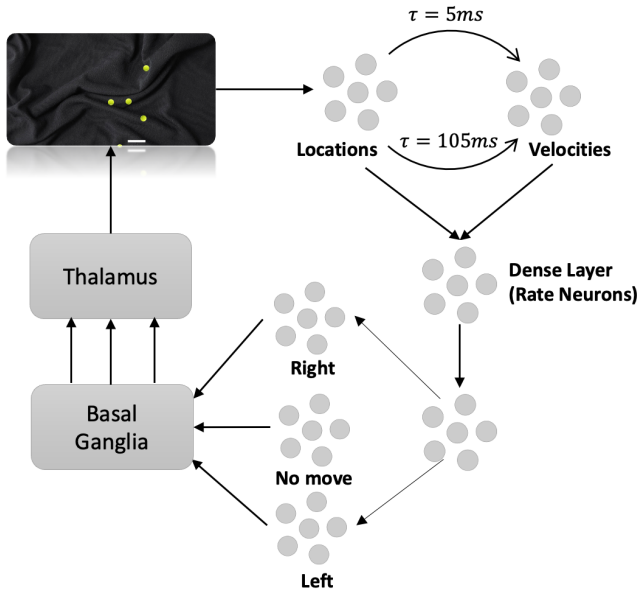

Figure 2: Detailed model architecture. The locations ensemble is connected to the velocities ensemble with a fast connection and a slow connection, and the difference between them is used to calculate the velocities. The information from both ensembles is connected to the dense layer, trained offline with SAC and the converted to spiking. The output of the dense layer undergoes an additional transform, resulting in a single number. This number is converted into left and right action utilities, which are joined by a threshold no move node, and collectively fed into the BG-Thalamus loop. The chosen action is sent back to the Unity environment.

Model Comparison Methodology

Model performance was compared to human minimum, median, and maximum scores to assess whether the neuromorphic architecture could reproduce the variability and capacity constraints observed in human executive function.

The ANN was trained without architectural constraints using the full observation vector provided by the game. In contrast, the fully spiking model was systematically constrained to test specific hypotheses about human performance variability. Object-tracking capacity was manipulated by varying the dimensionality of the attention and visual working memory (VWM) ensembles, corresponding to the number of balls encoded (two dimensions per ball for position and velocity). In addition, the number of neurons in each ensemble was varied to examine the effect of neural resources on performance.

After characterizing the effects of multi-object tracking limitations, two additional neural constraints were examined. First, perceptual noise was introduced into the attention ensemble using pink (1/f) Gaussian noise, consistent with known properties of human perceptual systems (Gilden et al., 1995). Second, model reaction time was manipulated by increasing synaptic time constants throughout the network within biologically plausible ranges, thereby varying the delay between sensory input and action selection. Reaction times were tested across a range matching established human variability (Welford et al., 1980).

Each model configuration was evaluated over 10 independent runs with randomized seeds to assess performance stability.

Results

Human Performance. Human participants (N = 26) displayed significant variability in game scores across the cohort, with total scores ranging from 88 to 191 points (Figure 3A), a population mean of 142 with standard deviation of 32, and a median of 156. There is a small negative correlation between age and score (Figure 3B). These findings are consistent with prior results showing significant differences in capacity when tracking multiple objects in dynamic environments and employing executive functions (Clayton et al., 2016; Green & Bavelier, 2006).

Machine-Optimized Performance. The ANN agent, trained using deep reinforcement learning, achieved stable superhuman performance on the task. Training performance increased steadily over five million steps and subsequently converged, with no further gains from additional training. The final average score of the Soft Actor-Critic agent (230) exceeded both the human median (155) and the highest observed human score (191).

These results demonstrate that the task admits a high performance ceiling under unconstrained optimization and establish a benchmark against which biologically constrained models can be compared. Consistent with prior work, the ANN's performance replicates the well-documented advantage of modern deep reinforcement learning agents over humans in complex, game-like environments (Badia et al., 2020; Berner et al., 2019).

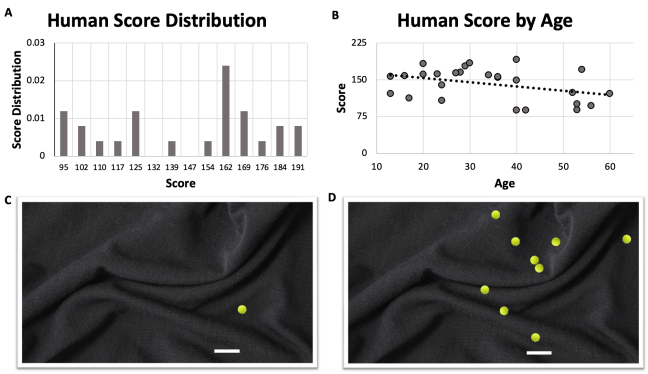

Figure 3: Human performance scores on the modified Pong game. (A) The distribution of human scores. The average reward is 142. The median is 155. (B) Human scores as a function of age, with a negative trendline. (C) The Pong game with one ball. (D) The Pong game with 9 balls.

Neuromorphic Model Performance. We isolated the effect of a single biological constraint by replacing the ANN's optimized output layer with the neuromorphic action selection module, comprising action-encoding ensembles and a basal-ganglia and thalamus circuit. This substitution alone reduced average performance to a score of 200, indicating a neuromorphic cost imposed by biologically plausible action selection dynamics. Importantly, this shift moved performance from a superhuman regime toward a more human-like range, prior to introducing perceptual or processing constraints.

We evaluated the full spiking neural network (SNN), which integrates neuromorphic perception, spiking processing layers, and basal ganglia action selection. To examine sources of performance variability, we systematically manipulated two parameters: (1) neural resources, defined by the number of neurons allocated to attention and visual working memory (VWM) ensembles, and (2) tracking capacity, defined by the number of balls the model encoded and tracked.

For low neuron counts, performance saturated at a small number of tracked objects, proportional to ensemble size. Beyond this saturation point, increasing tracking capacity failed to improve performance and often degraded it (Figure 4B). Two configurations were particularly informative: models with 200 and 2000 neurons per ensemble. In both cases, tracking a single ball yielded scores around 90, comparable to the human minimum. Performance increased monotonically with tracking capacity. In the 200-neuron model, scores saturated at four to five balls, yielding average scores comparable to the human median. In contrast, the 2000-neuron model continued to benefit from increased tracking capacity, peaking at nine balls with scores matching the highest-performing humans (Figure 4A).

To further probe mechanistic sources of variability, we examined perceptual noise in the strongest configuration (2000 neurons, nine balls). Adding pink Gaussian noise to the attention ensemble produced a monotonic decline in performance (Figure 4C). Low noise levels yielded peak

human-like scores, moderate noise reduced performance to the human median, and high noise collapsed performance to the human minimum.

Finally, we manipulated model reaction time by varying synaptic time constants, producing latencies between 100 and 300 ms. Performance declined monotonically with increasing reaction time (Figure 4D). Short reaction times matched peak human performance, intermediate reaction times aligned with the human median, and long reaction times reproduced the lowest human scores.

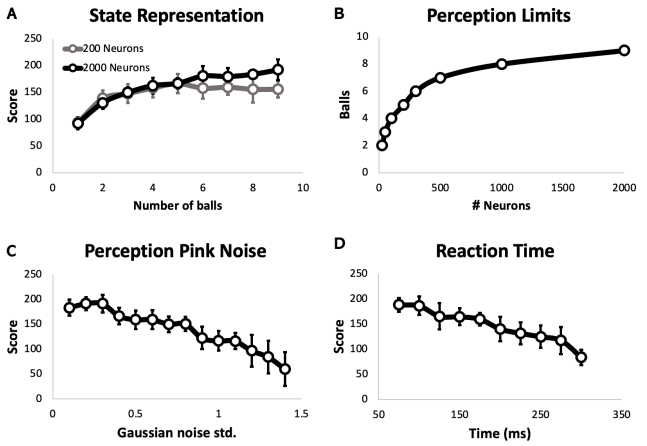

Figure 4: Neuromorphic model performance, averaged over 10 runs.

Discussion

This work proposes a neurally mechanistic account of the limits and variability of human executive function in a dynamic task. By combining human behavioral data with two computational models, one optimized for performance and the other constrained by biological plausibility, we reproduced key characteristics of human behavior in multi object tracking, including pronounced interindividual variability and capacity limits.

Human performance in the modified Pong task exhibited a broad distribution of scores, consistent with prior evidence that executive functions are multifaceted and vary substantially across individuals (Friedman et al., 2008). An unconstrained ANN trained with deep reinforcement learning achieved superhuman performance, demonstrating that the task itself does not impose a low ceiling. In contrast, a biologically inspired SNN exhibited performance that depended parametrically on neural constraints, including the number of neurons allocated to tracking, perceptual noise, and reaction time. Varying these parameters generated a range of scores closely matching the human distribution.

This result provides a neurally grounded instantiation of the flexible resource theory of attention. While this theory posits that attentional capacity reflects the division of a finite resource (Ma et al., 2014), the nature of this resource has remained abstract. In our model, the number of neurons allocated to object tracking directly constrained performance: for a given neural resource level, performance saturated at a specific number of tracked objects, and attempting to track additional objects led to degraded representations and lower scores. This suggests that neural population size is a plausible candidate mechanism underlying flexible attentional resources.

Our model corresponds with human performance. When restricted to tracking a single object, the model reproduced the lowest human scores, consistent with participants adopting minimal strategies. With approximately 200 neurons in the perception module, performance saturated at four to five objects, closely matching typical human capacity limits and the median of the human cohort (Fougnie & Marois, 2006). Increasing the neural resources to 2000 neurons enabled tracking of more objects and produced scores comparable to the highest-performing humans. These results support the view that individual differences in executive performance can arise from differences in both strategy and available neural resources, shaped by training and genetic factors (Astle & Scerif, 2011; Bays et al., 2009).

In addition to capacity limits, human variability has been attributed to perceptual noise and reaction time. Our model reproduced the full range of human scores by varying pink noise levels and reaction times within empirically plausible ranges. These effects emerged within the same spiking architecture that exhibited tracking limits, demonstrating that capacity, noise, and reaction time arise from shared neural constraints rather than independent factors.

Interestingly, while the performance-optimized ANN far exceeded human performance, replacing its output stage with a biologically plausible basal ganglia–thalamic circuit reduced its performance. This result might suggest that biological circuits are not optimized for maximal task performance, but for robustness, flexibility, and generalized action selection (Palmeri et al., 2017).

Several limitations point to future directions. The model employs a simplified representation of visual working memory, without prioritization among stored items. Incorporating uneven resource allocation could further refine performance predictions (Ma et al., 2014). In addition, synaptic weights were analytically derived rather than learned, and the perception module relied on ANN-to-SNN conversion. Future work chould integrate biologically plausible learning rules and fully spiking reinforcement learning methods. Finally, while the modified Pong task captures key aspects of dynamic executive control, extending the model to more complex, multi-task environments will be essential for assessing generalizability.

This study provides a neurally mechanistic account of the limits and variability of human executive function. A biologically plausible spiking neural network reproduced the range of human performance through variation in tracking capacity, noise, and reaction time, highlighting the role of biological constraints in shaping human-like behavior.

Acknowledgments

The authors would like to thank the members of the Neuro-Biomorphic Engineering Lab at the Open University of Israel for the insightful discussions

References

Abbott, L. F., & Dayan, P. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press.

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, *9*, 357–381. https://doi.org/10.1146/annurev.ne.09.030186.002041

Alvarez, G. A., & Franconeri, S. L. (2007). How many objects can you track?: Evidence for a resource-limited attentive tracking mechanism. *Journal of Vision*, *7*(13), 14. https://doi.org/10.1167/7.13.14

Astle, D. E., & Scerif, G. (2011). Interactions between attention and visual short-term memory (VSTM): What can be learnt from individual and developmental differences? *Neuropsychologia*, *49*(6), 1435–1445. https://doi.org/10.1016/j.neuropsychologia.2010.12.001

Badia, A. P., Piot, B., Kapturowski, S., Sprechmann, P., Vitvitskyi, A., Guo, D., & Blundell, C. (2020). *Agent57: Outperforming the Atari Human Benchmark* (No. arXiv:2003.13350). arXiv. https://doi.org/10.48550/arXiv.2003.13350

Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), 7.1-11. https://doi.org/10.1167/9.10.7

Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science (New York, N.Y.)*, *321*(5890), 851–854. https://doi.org/10.1126/science.1158023

Bekolay, T., Bergstra, J., Hunsberger, E., DeWolf, T., Stewart, T. C., Rasmussen, D., Choo, X., Voelker, A. R., & Eliasmith, C. (2014). Nengo: A Python tool for building large-scale functional brain models. *Frontiers in Neuroinformatics*, *7*, 48. https://doi.org/10.3389/fninf.2013.00048

Bellemare, M. G., Naddaf, Y., Veness, J., & Bowling, M. (2013). The Arcade Learning Environment: An Evaluation Platform for General Agents. *Journal of Artificial Intelligence Research*, *47*, 253–279. https://doi.org/10.1613/jair.3912

Berner, C., Brockman, G., Chan, B., Cheung, V., Dębiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., Pinto, H. P. d O., Raiman, J., Salimans, T., Schlatter, J., … Zhang, S. (2019). *Dota 2 with Large Scale Deep Reinforcement Learning* (No. arXiv:1912.06680). arXiv. https://doi.org/10.48550/arXiv.1912.06680

Blakowski, I., Zendrikov, D., Indiveri, G., & Capone, C. (2025). Neuromorphic dreaming as a pathway to efficient learning in artificial agents. *Neuromorphic Computing and Engineering*, *5*(4), 044005. https://doi.org/10.1088/2634-4386/ae0a77

Clayton, K. K., Swaminathan, J., Yazdanbakhsh, A., Zuk, J., Patel, A. D., & Kidd, G. (2016). Executive Function, Visual Attention and the Cocktail Party Problem in Musicians and Non-Musicians. *PLoS ONE*, *11*(7), e0157638. https://doi.org/10.1371/journal.pone.0157638

Cohen-Duwek, H., Slovin, H., & Ezra Tsur, E. (2022). Computational modeling of color perception with biologically plausible spiking neural networks. *PLoS Computational Biology*, *18*(10), e1010648. https://doi.org/10.1371/journal.pcbi.1010648

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *The Behavioral and Brain Sciences*, *24*(1), 87–114; discussion 114-185. https://doi.org/10.1017/s0140525x01003922

Cowan, N. (2010). The Magical Mystery Four: How is Working Memory Capacity Limited, and Why? *Current Directions in Psychological Science*, *19*(1), 51–57. https://doi.org/10.1177/0963721409359277

Cross, L., Cockburn, J., Yue, Y., & O'Doherty, J. P. (2021). Using deep reinforcement learning to reveal how the brain encodes abstract state-space representations in high-dimensional environments. *Neuron*, *109*(4), 724-738.e7. https://doi.org/10.1016/j.neuron.2020.11.021

Destexhe, A., Mainen, Z., & Sejnowski, T. J. (2003). Kinetic models for synaptic interactions. In *Methods in Neuronal Modeling: From Ions to Networks* (2nd ed.). MIT Press. https://www.researchgate.net/publication/268359503_Kinetic_models_for_synaptic_interactions

Diamond, A. (2013). Executive Functions. *Annual Review of Psychology*, *64*, 135–168. https://doi.org/10.1146/annurev-psych-113011-143750

Eliasmith, C., & Anderson, C. H. (2003). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. MIT Press.

Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., & Rasmussen, D. (2012). A large-scale model of the functioning brain. *Science (New York, N.Y.)*, *338*(6111), 1202–1205. https://doi.org/10.1126/science.1225266

Fan, J. (2023). *A Review for Deep Reinforcement Learning in Atari:Benchmarks, Challenges, and Solutions* (No. arXiv:2112.04145). arXiv. https://doi.org/10.48550/arXiv.2112.04145

Fougnie, D., & Marois, R. (2006). Distinct Capacity Limits for Attention and Working Memory: Evidence From Attentive Tracking and Visual Working Memory Paradigms. *Psychological Science*, *17*(6), 526–534. https://doi.org/10.1111/j.1467-9280.2006.01739.x

Friedman, N. P., Miyake, A., Young, S. E., DeFries, J. C., Corley, R. P., & Hewitt, J. K. (2008). Individual Differences in Executive Functions Are Almost Entirely Genetic in Origin. *Journal of Experimental Psychology. General*, *137*(2), 201–225. https://doi.org/10.1037/0096-3445.137.2.201

Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal Population Coding of Movement Direction. *Science*, *233*(4771), 1416–1419. https://doi.org/10.1126/science.3749885

Gilden, D. L., Thornton, T., & Mallon, M. W. (1995). 1/f Noise in Human Cognition. *Science*, *267*(5205), 1837–1839. https://doi.org/10.1126/science.7892611

Green, C. S., & Bavelier, D. (2006). Enumeration versus multiple object tracking: The case of action video game players. *Cognition*, *101*(1), 217–245. https://doi.org/10.1016/j.cognition.2005.10.004

Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). *Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor* (No. arXiv:1801.01290). arXiv. https://doi.org/10.48550/arXiv.1801.01290

Hunsberger, E., & Eliasmith, C. (2015). Spiking Deep Networks with LIF Neurons. *ArXiv*. https://www.semanticscholar.org/paper/Spiking-Deep-Networks-with-LIF-Neurons-Hunsberger-Eliasmith/0a76cdb10b244d90770888f72ee46ca3ea1d918d

Juliani, A., Berges, V.-P., Teng, E., Cohen, A., Harper, J., Elion, C., Goy, C., Gao, Y., Henry, H., Mattar, M., & Lange, D. (2020). *Unity: A General Platform for Intelligent Agents* (No. arXiv:1809.02627). arXiv. https://doi.org/10.48550/arXiv.1809.02627

Kagan, B. J., Kitchen, A. C., Tran, N. T., Habibollahi, F., Khajehnejad, M., Parker, B. J., Bhat, A., Rollo, B., Razi, A., & Friston, K. J. (2022). *In vitro* neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, *110*(23), 3952-3969.e8. https://doi.org/10.1016/j.neuron.2022.09.001

Kang, M.-S., Hong, S. W., Blake, R., & Woodman, G. F. (2011). Visual working memory contaminates perception. *Psychonomic Bulletin & Review*, *18*(5), 860–869. https://doi.org/10.3758/s13423-011-0126-5

Konda, V., & Tsitsiklis, J. (1999). Actor-Critic Algorithms. *Advances in Neural Information Processing Systems*, *12*. https://papers.nips.cc/paper_files/paper/1999/hash/6449f44a102fde848669bdd9eb6b76fa-Abstract.html

Kriegeskorte, N., & Douglas, P. K. (2018). Cognitive computational neuroscience. *Nature Neuroscience*, *21*(9), 1148–1160. https://doi.org/10.1038/s41593-018-0210-5

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, *17*(3), 347–356. https://doi.org/10.1038/nn.3655

Maass, W. (1997). Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, *10*(9), 1659–1671. https://doi.org/10.1016/S0893-6080(97)00011-7

Michel, M., & Geisler, W. S. (2011). Intrinsic position uncertainty explains detection and localization performance in peripheral vision. *Journal of Vision*, *11*(1), 18. https://doi.org/10.1167/11.1.18

Oryshchuk, A., Sourmpis, C., Weverbergh, J., Asri, R., Esmaeili, V., Modirshanechi, A., Gerstner, W., Petersen, C. C. H., & Crochet, S. (2024). Distributed and specific encoding of sensory, motor, and decision information in the mouse neocortex during goal-directed behavior. *Cell Reports*, *43*(1), 113618. https://doi.org/10.1016/j.celrep.2023.113618

Palmeri, T. J., Love, B. C., & Turner, B. M. (2017). Model-based cognitive neuroscience. *Journal of Mathematical Psychology*, *76*(Pt B), 59–64. https://doi.org/10.1016/j.jmp.2016.10.010

Potjans, W., Morrison, A., & Diesmann, M. (2009). A spiking neural network model of an actor-critic learning agent. *Neural Computation*, *21*(2), 301–339. https://doi.org/10.1162/neco.2008.08-07-593

Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, *3*(3), 179–197. https://doi.org/10.1163/156856888x00122

Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the selection problem? *Neuroscience*, *89*(4), 1009–1023. https://doi.org/10.1016/s0306-4522(98)00319-4

Roussy, M., Mendoza-Halliday, D., & Martinez-Trujillo, J. C. (2021). Neural Substrates of Visual Perception and Working Memory: Two Sides of the Same Coin or Two Different Coins? *Frontiers in Neural Circuits*, *15*. https://doi.org/10.3389/fncir.2021.764177

Stewart, T., Choo, X., & Eliasmith, C. (2010). *Dynamic Behaviour of a Spiking Model of Action Selection in the Basal Ganglia Neural Structure*. https://www.semanticscholar.org/paper/Dynamic-Behaviour-of-a-Spiking-Model-of-Action-in-Stewart-Choo/20188e51bab9afb5b1c42d0bc3b21cad61d48bd7

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning, second edition: An Introduction*. Bradford Books.

Takahashi, N., Ono, T., Omori, Y., Iizumi, M., Kato, H., Kasuno, S., Persing, B., & Tsuchiya, K. J. (2024). Assessment of executive functions using a 3D-video game in children and adolescents with ADHD. *Frontiers in Psychiatry*, *15*. https://doi.org/10.3389/fpsyt.2024.1407703

Tsur, E. E. (2021). *Neuromorphic Engineering: The Scientist's, Algorithms Designer's and Computer Architect's Perspectives on Brain-Inspired Computing*. CRC Press. https://www.taylorfrancis.com/books/mono/10.1201/9781003143499/neuromorphic-engineering-elishai-ezra-tsur

Tsur, E. E., & Rivlin-Etzion, M. (2020). Neuromorphic implementation of motion detection using oscillation interference. *Neurocomputing*, *374*, 54–63. https://doi.org/10.1016/j.neucom.2019.09.072

Welford, W. T., Brebner, J. M. T., & Kirby, N. (1980). *Reaction Times*. Stanford University.

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), 233–235.