Allostatic Control of Persistent States in Spiking Neural Networks for Perception and COMPUTATION

Anonymous authors

006

008 009 010

011

013

014

015

016

017

018

019

021

023

025

026

027

028

029

031

032

035

Paper under double-blind review

ABSTRACT

We introduce a novel model for updating perceptual beliefs about the environment by extending the concept of Allostasis to the control of internal representations. Allostasis is a fundamental regulatory mechanism observed in animal physiology that orchestrates responses to maintain a dynamic equilibrium in bodily needs and internal states. In this paper, we focus on an application in numerical cognition, where a bump of activity in an attractor network is used as a spatial-numerical representation. While existing neural networks can maintain persistent states, to date, there is no unified framework for dynamically controlling spatial changes in neuronal activity in response to environmental changes. To address this, we couple a well-known allostatic microcircuit, the Hammel model, with a ring attractor, resulting in a Spiking Neural Network architecture that can modulate the location of the bump as a function of some reference input. This localized activity in turn is used as a perceptual belief in a simulated subitization task – a quick enumeration process without counting. We provide a general procedure to fine-tune the model and demonstrate the successful control of the bump location. We also study the response time in the model with respect to changes in parameters and compare it with biological data. Finally, we analyze the dynamics of the network to understand the selectivity and specificity of different neurons to distinct categories present in the input. The results of this paper, particularly the mechanism for moving persistent states, are not limited to numerical cognition but can be applied to a wide range of tasks involving similar representations

034 1 INTRODUCTION

Neural networks in the brain are structured networks with a space-like structure, topologically analogous to a metric space (i.e. Euclidean) (Gerstner et al. (2014)). The use of these organized structure 037 are found not only in concrete representations, such as head direction (Pisokas et al. (2020)), orientation and compass for foraging (Stentiford et al. (2024)), and place and grid cells (Whelan et al. (2022)), but also in more abstract topographic maps that represent emotions, relationships, and, as 040 we show in this paper, numbers (Leslie et al. (2008)). Moreover, these structures extend beyond 041 biological systems; they are increasingly being employed in artificial embodied agents, demon-042 strated through tasks such as numerical cognition through embodied learning mechanisms and on 043 development of grounding transfer from concrete concept of sensorimotor experiences to abstract 044 conceptual knowledge (Alessandro & Angelo (2021)).

A key property of these networks is their capacity to hold persistent states, or bumps of activity (Gerstner et al. (2014)), which are associated with the representation of information in all of the mentioned systems of concrete and abstract representations (Hopfield (2015)). In this paper, we explore one such key network, the ring attractor model, which uses a bump of activity that is localized to dynamically control spatial changes in neuronal activity (i.e., activities influenced by changes in surrounding contextual features) (Pisokas et al. (2020)). Although these networks can represent to track changes in the environment, this raises the question of how we can exert arbitrary control over the position of bump activity in a structured spiking neural network. We infer that a structured network (i.e., ring attractor models, place cell and grid cells, Kohonen maps) would more closely resemble the brain's functionality than continuous neural field models, for which results in these 054 continuous neural field models already exists (Gerstner et al. (2014)). Moreover, we pose a question 055 on whether these bumps can represent abstract information that does not have explicit dynamics. 056

To approach the first question, we propose the use of self-regulation as a mechanism for controlling 057 the bump activity towards a desired state. The self-regulation mechanism here considers a neuronal model of set-point thermo-regulation, i.e., Hammel Model (Boulant (2006)), where this acts as a negative feedback control system for the bump activity. Considering that self-regulation can 060 achieve arbitrary control of manifold physiological variables, these variables have similar properties 061 to the bump activity (i.e., there is always a "place" where the bumps are supposed to be). An im-062 portant aspect of self-regulation is allostasis (Sterling (2012)), which extends the simpler concept 063 of homeostasis (i.e. the maintenance of physiological set-points) to cover adaptive, predictive and 064 competitive aspects of self-regulation. Thus, we further extend the usage of allostasis to tackle perception and computation in structured neural networks by demonstrating the predictive behaviors 065 towards its desired state based on numerical cognition tasks. 066

067 Thus, we introduce subitization to approach the second question as a predictive task over numer-068 ical capability. We use subitizing as a model system to explore abstract properties with neither 069 intrinsic nor explicit dynamics, yet with a putative metric representation: the number line (Leslie et al. (2008)). Subitization is a special mechanism capable of recognizing quantity strictly through 071 perception without the need for counting (Dehaene & Cohen (1994)) and it can be characterized as a form of categorization (Benoit & Henri Lehalle (2004)). This requires the distinction of ele-072 ments and their mapping to specific internal representations; in this case, bump activity is mapped 073 to numerical representations. Since subitization tasks are associated with the prediction of discrete 074 number representations, this form of representations can be modeled as an accumulator model with 075 bi-directional mappings (Gallistel & Gelman (2000)). Thus, we can consider numerosity tasks, such 076 as subitization, to exemplify the spatial-temporal control of these mappings. While there are other 077 models that encode numerosity tasks into a single neuron network (Rapp & Martin Paul Nawrot 078 (2020)), we use our allostatic network approach to achieve results that can match behavioral and 079 neuronal findings.





1.1 MODEL DERIVATION

081

082

090

092 093

094

096

099 100 101

109	Table 1: Default neuron parameters from NEST simulator are used for the network. The following
110	parameters are initially based on (Laing & Chow (2001)), where a parameter search is done to find
111	the regime for the ring attractor drift to be linear to the Hammel input. In the case of changes made,
112	these parameters are mentioned below and abbreviations used for parts of the model can be referred
112	from Figure 1

113	Name	Description	Value
114	V_{th}	Threshold Function	0.8mV
115	$\tau^1_{sun\ ex}$	Excitatory Synaptic Alpha Function(Hammel Model, S)	75ms
116	τ^2_{super}	Excitatory Synaptic Alpha Function (Ring Attractor)	27ms
117	τ_{super}^{3}	Excitatory Synaptic Alpha Function (HGM, LGM)	1000ms
118	τ_{super}^4	Excitatory Synaptic Alpha Function (Slow Inhibitory Neuron,	100ms
119	syn_cx	Number Discrimination Neuron)	
120	$\tau_{sun\ ex}^5$	Excitatory Synaptic Alpha Function (I)	1ms
121	$ au_{syn_in}$	Inhibitory Synaptic Alpha Function	50ms
122	I_e	Input Current (Default)	$0.45 \times 10^6 \mathrm{pA}$
123	I_e^1	Input Current (HGM, LGM, Slow Inhibitory Neuron,	0pA
124	C	Number Discrimination Neuron)	-
125	$ au_m$	Membrane Time Constant (Default)	7.04ms
126	$ au_m^1$	Membrane Time Constant (HGM, LGM)	4.2ms
127	$ au_m^2$	Membrane Time Constant (I)	7.06ms
128	C_m	Membrane Capacitance (Default)	$3.96 \times 10^6 \mathrm{pF}$
129	C_m^1	Membrane Capacitance (HGM, LGM, I)	$0.8 imes 10^6 \mathrm{pF}$

108

The Allostatic Controller of Persistent States architecture (AlloNet) starts with a group of neurons 132 that function as afferent neurons, providing input for the downstream tasks to be done by the external 133 perception neuron (figure 1A). These inputs are used by the Hammel model to self-regulate internal 134 beliefs, representing the interaction with the environment. At its core, we implement a spatially 135 structured network, i.e. the ring attractor network, with N neurons that are fine-tuned to maintain 136 persistent states or bumps. These persistent states are defined to be the internal beliefs that will be 137 controlled by the Hammel model. Each neuron in the model is an Integrate and fire (IAF) neuron, 138 and all the synapses are based on alpha synapses (Gerstner et al. (2014)) with decay constant τ_{ex}^k , 139 where k varies based on the specific sub-networks.

140 The synaptic weights of these neurons are defined by a symmetry of local excitation and surrounding 141 inhibition to ensure the stability of the persistent states. These synaptic weights, w, can be defined 142 by the following equation: 143

$$w_{ij} = \frac{\sigma_2 e^{-\frac{d_{ij}^2}{2\sigma_1^2}} - \sigma_1 e^{-\frac{d_{ij}^2}{2\sigma_2^2}}}{\sigma_2 - \sigma_1}$$
(1)

145 146 147

148

149

144

Here, i represents the pre-synaptic neuron and j represents the post-synaptic neuron. The terms

 $e^{-\frac{d_{ij}^2}{2\sigma_1^2}}$ and $e^{-\frac{d_{ij}^2}{2\sigma_2^2}}$ are Gaussian functions that define the spread of the inhibition and excitation 150 151 within the network, where the standard deviation σ_2 and σ_1 act as parameters controlling the width of the two Gaussian functions. The parameters d_{ij} here is a distance-dependent coupling between 152 the pre-synaptic and post-synaptic neuron, where we can consider as d(|i-j|). To scale the weights 153 for larger or smaller network populations, we define $\sigma_n = \frac{\delta_n 100}{N}$, where $n \in \{1, 2\}$ and N is the population size. The model used in our experiment consists of (N = 100) neurons with $\sigma_2 = 5$ 154 155 and $\sigma_1 = 10$ being used for tuning the parameters of the network. Thus, the δ_n parameter is the 156 standard deviation of the Gaussian function that will be scaled based on this population and standard 157 deviation parameters. The tuned parameters of the ring attractor network are provided in table 1. 158

To control the position, we use a homeostatic circuit inspired by the Hammel model of autonomic 159 nervous system for temperature regulation in mammals (Boulant (2006)). This model is composed 160 of four neurons (figure 1B). In the original model, an insensitive neuron (I) receives input for the 161 model and a sensitive neuron (S) function as a set-regulatory point. The excitation-inhibition balance

162 of these two neurons is processed through two downstream neurons. These neurons control opposite 163 effectors that aim to either generate or lose heat. This depends on a reference signal relayed to some 164 sensitive neurons by heat receptors in the skin or internal tissues. We adapt this concept to align 165 internal representations, i.e. the localized bump activity representing the system's predicted belief, 166 toward the external inputs in our model. In our adaptation of the model, the reference signal that functions similarly to the insensitive neuron corresponds towards system's predicted belief, and the 167 sensitive neuron that are typically a set-regulatory point can be dynamically adjusted to responses 168 from the environment. This meant that the effectors are the system's belief requiring a mechanism to shift its perception. Thus, we introduce a method that would allow the Hammel Model to shift 170 the localized activity of the ring attractor and be fed back into the network to form a closed control 171 loop. 172

It is known that under a symmetric kernel with local excitation and surrounding inhibition, the 173 ring attractor network can hold persistent activities in a localized region (Gerstner et al. (2014)). 174 These activities can be directed in a particular direction within the ring attractor by breaking the 175 symmetry of activities in the network. Typically, these are achieved through offsets implemented by 176 an asymmetric connection (Khona & Fiete (2022); Zhang (1996)). Thus, we hypothesize that this 177 control on shifting the position of the bump can be maintained by breaking the symmetry through 178 gain modulation, which will be explained in the next section. This mechanism is in charge of an 179 outer layer of two population of neurons connected to the ring attractor as shown in figure (1C). The homeostatic effector neurons of the Hammel model are connected to the Left Gain Modulation 181 (LGM) and (RGM) neurons to simulate the process of generating or losing heat, respectively.

182 Finally, a readout network is defined to feed back into the homeostasis model as a reference signal, 183 where these inputs are based upon the position of the bump in the ring attractor (figure 1D). This 184 circuit is inspired from models of familiarity discrimination that occur in the Perirhinal Cortex for 185 processing information as a downstream task (Bogacz R (2001)). This model is used for creating 186 a reference signal for representation perception neurons (figure 1B) from a cumulation of ring at-187 tractor neurons' responses. This usually involves three layer of neurons: representation neurons, 188 implemented by the neurons in the ring attractor, familiarity discrimination neurons, referred to as 189 the number discrimination neurons in our model; and decision neurons, implemented by the repre-190 sentation perception neuron in the Hammel Model (sensitive neurons) (figure 1D). These decision neurons project downstream specific rate encodings based on representations from the ring attractor. 191 The external output that drives the changing threshold in the Hammel model as a function of the 192 input works as an accumulation of responses from the environment (figure 1D). A use case for these 193 mechanisms in our model will be presented later. 194

195 196

197

2 CONTROLLING THE BUMP

The allostatic control of the bump in the ring attractor is achieved in two stages, which can be fine-tuned depending on the application. In the first stage, we determine the parameters of the homeostasis model such that the effector neurons vary (almost) linearly as function of the difference in input rate between sensitive and insensitive neurons (figure 2A). In this figure, we have isolated the homeostasis model and replaced the inputs with $I_{syn} = I_{poisson} + I_{bg}$, where $I_{poisson}$ represents the simulated Poisson spike train with a given rate, and I_{bg} is the background noise. This background noise helps linearize the firing rate of inputs to the homeostasis model.

In the second stage, we tune the gain modulation neurons so that they can shift the position of the bump (figure 2B). To achieve this, we first note that the architecture shown in figure 1C implies that the input to the *i*th LGM neuron has two terms:

208 209

212

213

$$I_{syn}^{i} = I_{ring}^{i-1} + I_{HL},$$
(2)

where I_{HL} is the output from the "Low" effector of the homeostasis model, and I_{ring}^{i-1} is the output from the ring attractor. Similarly, the input *i*th HGM neuron is given by

$$I_{syn}^{i} = I_{ring}^{i+1} + I_{HH}, (3)$$

where I_{HH} is the output from the "High" effector of the homeostasis model. Since *i* represents the neuron index of *N* in the ring attractor, we consider the *i* for I_{HL} and IHH to their respective positions in the ring attractor.



Figure 2: Overview the fine-tuning stages for the model. Each time-step used by the simulator is 230 equivalent to an interval of 1ms. A. Heat map for the firing rate in ratio to the two inputs to the 231 Hammel model, simulated by a Poisson generator. Left. Activity (Firing rate) in lower neuron 232 for different input frequencies. Right. Activity of higher neuron. The model is simulated for 10000 233 time-steps for each combination of input frequencies. B. Setup for the fine-tuning process of the gain 234 modulation neurons as described in the text. A similar setup from diagram A. Poisson generators 235 are then used for simulating the bump of the ring attractor and the Hammel model effector output. 236 The colors represent firing rates in response to different combinations of firing. C. Response of the 237 bump for two different frequencies of the external input (50000 time-steps). Left. Raster plot. Right. 238 Average firing rate of the readout neurons. The colors represent the ranges for different numbers (see 239 below).

246

Each of these neurons, in turn, inhibits the next (or previous) ring attractor neuron. Consequently, whenever the ring and the Hammel effector neurons are co-active, they will unbalance the activity in the ring attractor. Moreover, we need to guarantee that when only one of the two (either the neuron in the ring attractor or the corresponding output of the homeostasis model) is active, the LGM or HGM neuron is close to the threshold and do not fire:

$$V_m(I_{ring}^{i\pm 1}) \approx V_m(I_{HH}) = V_{th} - h < V(I_{ring}^{i\pm 1} + I_{HH}),$$
(4)

where V_m is the membrane potential of the LGM or HGM neurons, V_{th} is the firing threshold and *h* is a small arbitrary constant (figure 2B). With these considerations, we can move the bump so that it matches internal and external representations (figure 2C). All the simulations are performed using NEST simulator (Gewaltig & Diesmann (2007)) in a standard Laptop computer (13th Gen Intel® CoreTM i7-13700HX, 16GB Ram, Nvidia RTX4060). The parameters of the Hammel and Gain modulation networks are given in table 1.

253 254 255

256

257

258

267

3 SUBITIZING AS AN ALLOSTATIC PROCESS: BEHAVIORAL RESPONSES

As an example, we introduce the AlloNet as a model of subitizing. Subitizing can be seen as a "number sense" that we define, in this paper, as an allostatic match between external inputs and internal representations (figure 3A).

To set up our model to subitize, we define the ring attractor as a putative number line representation (Hamdan & Gunderson (2017); Revkin (2008)). The input to the network is represented as four Poisson spike train generators which feeds onto the insensitive neurons (here renamed as external perception neurons) and tuned to generate firing rates in the range [0 - 200]Hz depending on the number of co-activated inputs. Note that this guarantees permutation invariance of the inputs (i.e., activation sequences $\{1, 0, 1, 0\}$ generate the same input as $\{0, 1, 0, 1\}$ and so on).

We define the number representations by the bijection:

 $\{[0Hz, 50Hz], [51Hz, 100Hz], [101Hz, 150Hz], [151Hz, 200Hz]\} \mapsto \{1, 2, 3, 4\}.$ (5)

As seen in figure 2C, once the input is given, the model adjusts its internal representation to match to the corresponding numerosity. Similarly, we set up four readout neurons to match the corresponding frequencies as those of the input.



Figure 3: Simulated behavioral responses for a subitizing task. A. We use the AlloNet model to define a homeostatic match between the internal representation of numbers and stimuli in the environment (dots). B. Reaction times in association to number representation defined as the time of first spike in the correct representation. Responses in S2 neurons are used for determining reaction time. The experiments are repeated for different synaptic time constants of the HGM and LGM. C. Error percentage based on the responses at different stopping time-steps in 20 trials (see text) **D.** Bar plot to define quality score for different excitatory synaptic time constants and different numbers. The different bar represents the number being represented (color-coded) and the groupings of these bars are reported based on their synaptic time constants.

302

303

292

293

295

296

297

298

3.1 THE ALLONET MODEL COMPARISON TO BEHAVIORALLY PLAUSIBLE REACTIONS TIMES AND ERROR RATES

To investigate the relevance of our model as a model of subitizing, we simulated two common subitizing experiments in humans. In the first experiment, we tested whether the network reproduces relevant reaction times as those observed in humans. To do so, we define reaction time as the time of the first spike from the readout neuron corresponding to the input numerosity. The results are shown in figure 3B. We parameterize the experiment further by changing the excitatory synaptic time constant of the gain modulation network.

310 The model reproduces behavioral responses in two significant ways and fails in one crucial aspect. 311 For longer synaptic time constants (1000ms), the reaction time increases as a function of numerosity 312 and also becomes more variable. However, with shorter synaptic time constants (900ms), the model 313 presents paradoxical behavior: smaller numerosities (numbers less than 3) are noticeably slower 314 than higher numerosities, while this behavior reverses for numerosities higher than 3. In the case of 315 human data, it takes less than half a second to perceive the presence of one, two, or three objects, but the speed and accuracy fall dramatically beyond this limit (Dehaene (2011)). This human data 316 showcases two aspects of the model, where longer synaptic time constants (1000ms) aligned to 317 human responses for numbers less than 3, while short synaptic time constants (900ms) aligned with 318 the sudden increase in reaction time for humans on numbers greater than 3. However, the model 319 fails on the aspects of the dynamics being much slower than human reaction times. 320

An additional experiment evaluated the error rate of the model (figure 3c). To do so, we fixed a specific time during the run iteration in which the model should report the perceived numerosity and tested the error rate for four different ranges of response timesteps. Here, we define error rate as (number of misfired timestep ÷ total timestep). For early response times, error rate increases with numerosity as observed in experiments (Dehaene & Cohen (1994)). However, for longer waiting times, responses become inconsistent. We investigate this in the following section.

Lastly, we design an experiment to evaluate the performance of the ring attractor network by analyzing the representation quality over time within defined spatial intervals. Specifically, this analysis investigates a measure on reaching the desired state quickly while remaining stable within that predefined desired region. To calculate the quality score, we use the centroid calculation of the network, P(t). These are then classified as right or wrong time based on whether the centroid is in the desired region of bump activity. Thus, we define the quality score as follows:

$$Quality = \frac{\text{Number of right times}}{\text{Number of right times} + \text{Number of wrong times}}$$
(6)

The analysis shows that longer synaptic time constant performs better. Thus, we will discuss how different synaptic time constants affect the persistent states in the following section.

337

4 SUBITIZING AS AN ALLOSTATIC PROCESS: NEURAL DYNAMICS

To understand the responses for different excitatory synaptic time constants, we investigate through 342 a heat-map that visualizes the activity of neurons across different time windows (figure 4). The 343 excitatory synaptic time constant may affect the reaction time of the network due to two possible 344 outcomes: the slow angular speed of the shifting cues in the ring attractor network or instability of 345 persistent states in the network. Although the first outcome can be directly inferred as time taken 346 to be inversely proportional to the angular speed (Seen in $\tau = 900$ in comparison to $\tau = 1000$ 347 (figure 4)), the second outcome deters that high angular speed can also affect the network to have 348 slow response time. In the case that the network has excessive angular speed (greater than 0.018) 349 rad/ms), this may diminish the ring attractor's capability to track towards the desired persistent state 350 and oscillate between different signals (Chen (2024)) (Seen in $\tau = 1250$ in comparison to $\tau = 1000$ 351 (figure 4)).

As a further investigation of the network, we start by studying the responses of a single neuron to different numerosities over different trials. As an example, we show neuron number 63 in the ring attractor which is tuned to respond to number 3 (figure 3A).

The response of a given neuron is composed of transients for higher numerosities (as the bump transits its neighborhood) and persistent activity for the numerosity it is tuned to. Note however that for later times the bump becomes unstable and starts a new wander due to the noise of in the network. This explains the inconsistent error rates for higher response times in figure 3C.

We also studied the gain field in the vicinity of neuron 63, composed of only transient responses for numerosities greater than or equal to the tuned numerosity (figure 5A). These two profiles of activity resemble the responses of the Entorhinal Cortex, the Parahipocampal Cortex and the Hippocampus in humans (Kutter et al. (2023)).

To understand a bit better the dynamics of the bump, we studied the position coordinate (Hopfield (2015)) of the bump and its speed. The position coordinate consistently represents the numerosity across trials (figure 5B) and the velocity shows evidence of a constant acceleration towards the corresponding numerosity. Furthermore, by investigating the scaling of the speed with the synaptic time constant we discover the reason for the difference in reaction times for slower time constants: the bump movement profile becomes shallower for lower numerosities due to lack of excitation from the input (figure 5C, dotted line).

371 372

373

5 CONCLUSIONS AND FUTURE WORK

We have proposed a model of allostatic control of persistent states in spiking neural networks, named
the AlloNet. Additionally, we propose it as a computation model of subitizing in animals, defining
the number sense as a homeostatic alignment of internal representations and external inputs. Our
model reproduces qualitatively important behavioral aspects of subitizing in humans as well as observed neural dynamics of the regions known to be involved in number perception.



Figure 4: Heat-map to plot the network persistent state over time. The y-axis of the plot is the number representation index, and the y-axis of the sub-plot is the neuron index of the ring attractor network. The x-axis of the plot is the column for different synaptic time constant runs and the x-axis of the sub-plot is the time of the run iteration. The white lines in the sub-plots define the different numerosities, i.e., [starting point,1,2,3,4] from bottom to top.

420

421

422

423

The model is able "track" a static property of the environment by dynamically adjusting an spacelike internal model to match a function of the input which, due to the simplicity of the input network and the properties of neural integration, is numerosity. That is, this model follows the premise of perception as active process (Parr et al. (2022)), moreover, that perception is self-regulatory in that the internal and external models need to find a match through the lens of a shared interface (Parr et al. (2022)). This interface is instantiated in the Hammel model.



Figure 5: Neural dynamics and dynamics of the bump. The color is encoded to the location of the bump activity (shown by the number represented on the left of the plots in figure 5A) A. Responses of a single neuron of the ring attractor (63) sensitive to number 3. Top. Raster plot for 20 trials of each numerosity. Bottom. Corresponding firing rates. Right. Responses of the gain field in the vicinity of the ring (± 10 neurons). B. Centroid of the bump computed as the expected value of position (neuron) with respect to the normalized firing rate. The spike trains are also plotted in the background for each numerosity. C. Speed of the bump smoothed with 5 samples sliding window for each numerosity. Dotted line is the corresponding step for lower synaptic time constant ($\tau_{ex} = 900ms$).

449

450

451

452

453

454

455

456

459

The Hammel model is involved in self-regulation. It works by minimizing the error between a given input and a fixed set value. By setting variability and abstraction to the threshold (as numerosity is not a one-dimensional signal), we allowed the model to handle more general representations. We think that this idea can be extended further to solve more general computational problems (Hopfield (2015)).

The self-regulatory aspect of the model implies that there is an autonomic response (internal or covert change) and a behavioral response (external or cover). This is analogous to the duality between subitizing or approximating, counting (Dehaene & Cohen (1994)), both involve different systems and one of them is behavioral (Pecyna et al. (2022)) while the latter is purely autonomic or spontaneous (Castaldi et al. (2021)) and is modality independent (Togoli & Arrighi (2021)).

470 The self-regulatory aspect of subitizing is thus the main contribution of this work. This change 471 of perspective has the potential to unify different systems that share the same properties of the 472 current model under a common view that could be considered as a neuromorphic version of active 473 inference (Parr et al. (2022)). Indeed, similar models have been described in insects for controlling 474 head direction (Pisokas et al. (2020)) which have inspired computational (Stentiford et al. (2024)) 475 and robotic models (Robinson et al. (2022)). However, in contrast with many existing models, our 476 current model does not include any sort of learning but adaptation, a property that is key not only for subitizing (Togoli & Arrighi (2021)) but for general perception and control as well. We think 477 learning will add degrees of freedom to the model like the possibility of building a full Approximate 478 Number System. 479

There are still many challenges to overcome, and they are the focus of current computational and theoretical work. One of them is the fact that the model has to be slow to be stable (i.e. the difference between the fast time scales of spiking and the slow dynamics of the bump) needs to be significant. Furthermore, the qualitative descriptions here need to be validated against actual data that not only approaches the difference in age when subitizing but, perhaps more importantly, the dysfunction of such systems. We think our overall approach is productive in investigating numbers and similar abstract senses in animals as well as in robots.

486 REFERENCES

488 489	Di Nuovo Alessandro and Cangelosi Angelo. Abstract concept learning in cognitive robots. <i>Current Robotics Reports</i> , 2, 2021.
490 491	Laurent Benoit and François Jouen Henri Lehalle. Do young children acquire number words through subitizing or counting? <i>Cognitive Development</i> , 19(3):291–307, 2004.
492 493 494	Giraud-Carrier C. Bogacz R, Brown MW. Model of familiarity discrimination in the perirhinal cortex. <i>J Comput Neurosci.</i> , 10(1):5–23, 2001.
495 496	Jack A Boulant. Neuronal basis of hammel's model for set-point thermoregulation. <i>Journal of Applied Physiology</i> , 100(4):1347–1354, 2006.
497 498 499	Elisa Castaldi, Antonella Pomè, Guido Marco Cicchini, David Burr, and Paola Binda. The pupil responds spontaneously to perceived numerosity. <i>Nature communications</i> , 12(1):5944, 2021.
500 501	Yani et al Chen. Synaptic ring attractor: A unified framework for attractor dynamics and multiple cues integration. <i>Heliyon</i> , 10, 2024.
502 503 504	Stanislas Dehaene. <i>The number sense: How the mind creates mathematics, Rev. and updated ed.</i> Oxford University Press, 2011.
505 506 507	Stanislas Dehaene and Laurent Cohen. Dissociable mechanisms of subitizing and counting: neuropsychological evidence from simultanagnosic patients. <i>Journal of Experimental Psychology: Human perception and performance</i> , 20(5):958, 1994.
508 509 510	C.R. Gallistel and Rochel Gelman. Non-verbal numerical cognition: from reals to integers. <i>Trends in Cognitive Sciences</i> , 4(2):59–65, 2000.
511 512	Wulfram Gerstner, Werner M Kistler, Richard Naud, and Liam Paninski. <i>Neuronal dynamics: From single neurons to networks and models of cognition</i> . Cambridge University Press, 2014.
513 514	Marc-Oliver Gewaltig and Markus Diesmann. Nest (neural simulation tool). <i>Scholarpedia</i> , 2(4): 1430, 2007.
516 517	Noora Hamdan and Elizabeth A Gunderson. The number line is a critical spatial-numerical repre- sentation: Evidence from a fraction intervention. <i>Developmental psychology</i> , 53, 2017.
518 519	John J Hopfield. Understanding emergent dynamics: Using a collective activity coordinate of a neural network to recognize time-varying patterns. <i>Neural Computation</i> , 27(10):2011–2038, 2015.
520 521 522	Mikail Khona and Ila R Fiete. Attractor and integrator networks in the brain. <i>Nature Reviews Neuroscience</i> , 23, 2022.
523 524 525	Esther F Kutter, Gert Dehnen, Valeri Borger, Rainer Surges, Florian Mormann, and Andreas Nieder. Distinct neuronal representation of small and large numbers in the human medial temporal lobe. <i>Nature human behaviour</i> , 7(11):1998–2007, 2023.
526 527 528	Carlo Laing and Carson Chow. Stationary bumps in networks of spiking neurons. <i>Neural computa-</i> <i>tion</i> , 13:1473–94, 08 2001. doi: 10.1162/089976601750264974.
529 530	Alan M Leslie, Rochel Gelman, and CR Gallistel. The generative basis of natural number concepts. <i>Trends in cognitive sciences</i> , 12(6):213–218, 2008.
531 532 533	Thomas Parr, Giovanni Pezzulo, and Karl J Friston. Active inference: the free energy principle in mind, brain, and behavior. MIT Press, 2022.
534 535	Leszek Pecyna, Angelo Cangelosi, and Alessandro Di Nuovo. A robot that counts like a child: a developmental model of counting and pointing. <i>Psychological Research</i> , 86(8):2495–2511, 2022.
536 537	Ioannis Pisokas, Stanley Heinze, and Barbara Webb. The head direction circuit of two insect species. <i>Elife</i> , 9:e53985, 2020.
538 539	Hannes Rapp and Merav Stern Martin Paul Nawrot. Numerical cognition based on precise counting with a single spiking neuron. <i>iScience</i> , 23(2), 2020.

- Susannah K. et al Revkin. Does subitizing reflect numerical estimation? *Psychological Science*, 19, 2008.
- Brian S Robinson, Raphael Norman-Tenazas, Martha Cervantes, Danilo Symonette, Erik C Johnson, Justin Joyce, Patricia K Rivlin, Grace M Hwang, Kechen Zhang, and William Gray-Roncal.
 Online learning for orientation estimation during translation in an insect ring attractor network. *Scientific reports*, 12(1):3210, 2022.
- 547 Rachael Stentiford, James C Knight, Thomas Nowotny, Andrew Philippides, and Paul Graham.
 548 Estimating orientation in natural scenes: A spiking neural network model of the insect central complex. *bioRxiv*, pp. 2024–02, 2024.
- Peter Sterling. Allostasis: a model of predictive regulation. *Physiology & behavior*, 106(1):5–15, 2012.
- Irene Togoli and Roberto Arrighi. Evidence for an a-modal number sense: Numerosity adaptation
 generalizes across visual, auditory, and tactile stimuli. *Frontiers in Human Neuroscience*, 15:
 713565, 2021.
 - Matthew T Whelan, Alejandro Jimenez-Rodriguez, Tony J Prescott, and Eleni Vasilaki. A robotic model of hippocampal reverse replay for reinforcement learning. *Bioinspiration & Biomimetics*, 18(1):015007, 2022.
- K Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *Journal of Neuroscience*, 16(6):2112–2126, 1996. ISSN 0270-6474. doi: 10.
 1523/JNEUROSCI.16-06-02112.1996. URL https://www.jneurosci.org/content/16/6/2112.