



OPEN ACCESS

EDITED BY

Sheri Marina Markose,
University of Essex, United Kingdom

REVIEWED BY

Dimitrije Marković,
Technical University Dresden, Germany
Tom Ziemke,
Linköping University, Sweden

*CORRESPONDENCE

Oscar Guerrero Rosado,
oscar.guerrerososado@donders.ru.nl
Paul F. M. J. Verschure,
paul.verschure@donders.ru.nl

SPECIALTY SECTION

This article was submitted to
Computational Intelligence in Robotics,
a section of the journal
Frontiers in Robotics and AI

RECEIVED 24 September 2022

ACCEPTED 17 November 2022

PUBLISHED 02 December 2022

CITATION

Rosado OG, Amil AF, Freire IT and
Verschure PFMJ (2022), Drive
competition underlies effective
allostatic orchestration.
Front. Robot. AI 9:1052998.
doi: 10.3389/frobt.2022.1052998

COPYRIGHT

© 2022 Rosado, Amil, Freire and
Verschure. This is an open-access
article distributed under the terms of the
[Creative Commons Attribution License
\(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or
reproduction in other forums is
permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original
publication in this journal is cited, in
accordance with accepted academic
practice. No use, distribution or
reproduction is permitted which does
not comply with these terms.

Drive competition underlies effective allostatic orchestration

Oscar Guerrero Rosado*, Adrian F. Amil, Ismael T. Freire and Paul F. M. J. Verschure*

Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, Netherlands

Living systems ensure their fitness by self-regulating. The optimal matching of their behavior to the opportunities and demands of the ever-changing natural environment is crucial for satisfying physiological and cognitive needs. Although homeostasis has explained how organisms maintain their internal states within a desirable range, the problem of orchestrating different homeostatic systems has not been fully explained yet. In the present paper, we argue that attractor dynamics emerge from the competitive relation of internal drives, resulting in the effective regulation of adaptive behaviors. To test this hypothesis, we develop a biologically-grounded attractor model of allostatic orchestration that is embedded into a synthetic agent. Results show that the resultant neural mass model allows the agent to reproduce the navigational patterns of a rodent in an open field. Moreover, when exploring the robustness of our model in a dynamically changing environment, the synthetic agent pursues the stability of the self, being its internal states dependent on environmental opportunities to satisfy its needs. Finally, we elaborate on the benefits of resetting the model's dynamics after drive-completion behaviors. Altogether, our studies suggest that the neural mass allostatic model adequately reproduces self-regulatory dynamics while overcoming the limitations of previous models.

KEYWORDS

self-regulation, homeostasis, allostatic control, attractor model, need-based behavior, control theory

Introduction

As 19th-century physiologists Claude Bernard and Ivan Pavlov proposed, living systems are generally characterized by their ability to self-regulate (Bernard, 1865; Pavlov, 1955). Through self-regulation, an organism ensures its fitness by adjusting its inner processes relative to external perturbations (Papies and Aarts, 2016), in turn, assisting self-maintenance (i.e., Autopoiesis) (Maturana and Varela, 1991). One aspect of self-regulation is homeostasis, which describes the process of maintaining an internal state within a desirable range as proposed by Cannon (Cannon, 1939). Once external perturbations produce a deviation from the desirable range, a homeostatic error arises, driving a proportional error-correcting response to restore balance in the system. This process would be later formalized as a feedback control loop by Norbert Wiener (Wiener, 1948), father of cybernetics. However, the homeostatic control of a

single need is insufficient to ensure fitness since living organisms need to maintain a rather large set of internal needs with dynamically varying priorities from electrolysis to temperature and oxygenation. To overcome this and other gaps of homeostasis in explaining self-regulation, the notion of allostasis aims to capture the stability of the integrated self rather than its parts.

Allostasis transcends the constancy imposed by homeostasis by aligning all organism's internal parameters with the environmental demands and opportunities (Sterling, 1988). In its broad conceptualization, allostasis targets stability by integrating many dynamic regulatory principles (Sterling, 2020). For this reason, we differentiate three complementary and coupled levels of allostasis. First, allostatic orchestration allows the agent to rapidly rank its internal needs based on priorities, urgencies, and opportunities. Second, predictive allostasis leverages environmental regularities to learn associations between external events and internal states, supporting anticipatory allostatic orchestration and future homeostatic risks. Finally, contextual allostasis benefits from goal-oriented learning mechanisms to enrich self-regulatory strategies with spatio-temporal information, in turn facilitating anticipatory behavioral strategies. Hence, we consider that in the mammalian brain allostatic regulation of action is organized in a multi-layered architecture following the Distributed Adaptive Control (DAC) theory (Verschure, 2012, 2016).

While nowadays, allostasis is increasing in popularity (2010: 2,584 citations *versus* 2021:12549 citations, resource: Scopus), most of the latest computational models on allostasis focus on exploring the advantages of its predictive component (Sterling, 2012). Paradoxically, some models return to the one-single-need problem of homeostasis (Tschantz et al., 2022). In addition, the concept of allostasis has not explained the computational mechanisms by which individuals achieve stability by orchestrating different homeostatic systems. The divergent modeling approaches adopted by the few computational studies addressing allostasis demonstrate the lack of consensus when determining the fundamental principles behind allostatic orchestration.

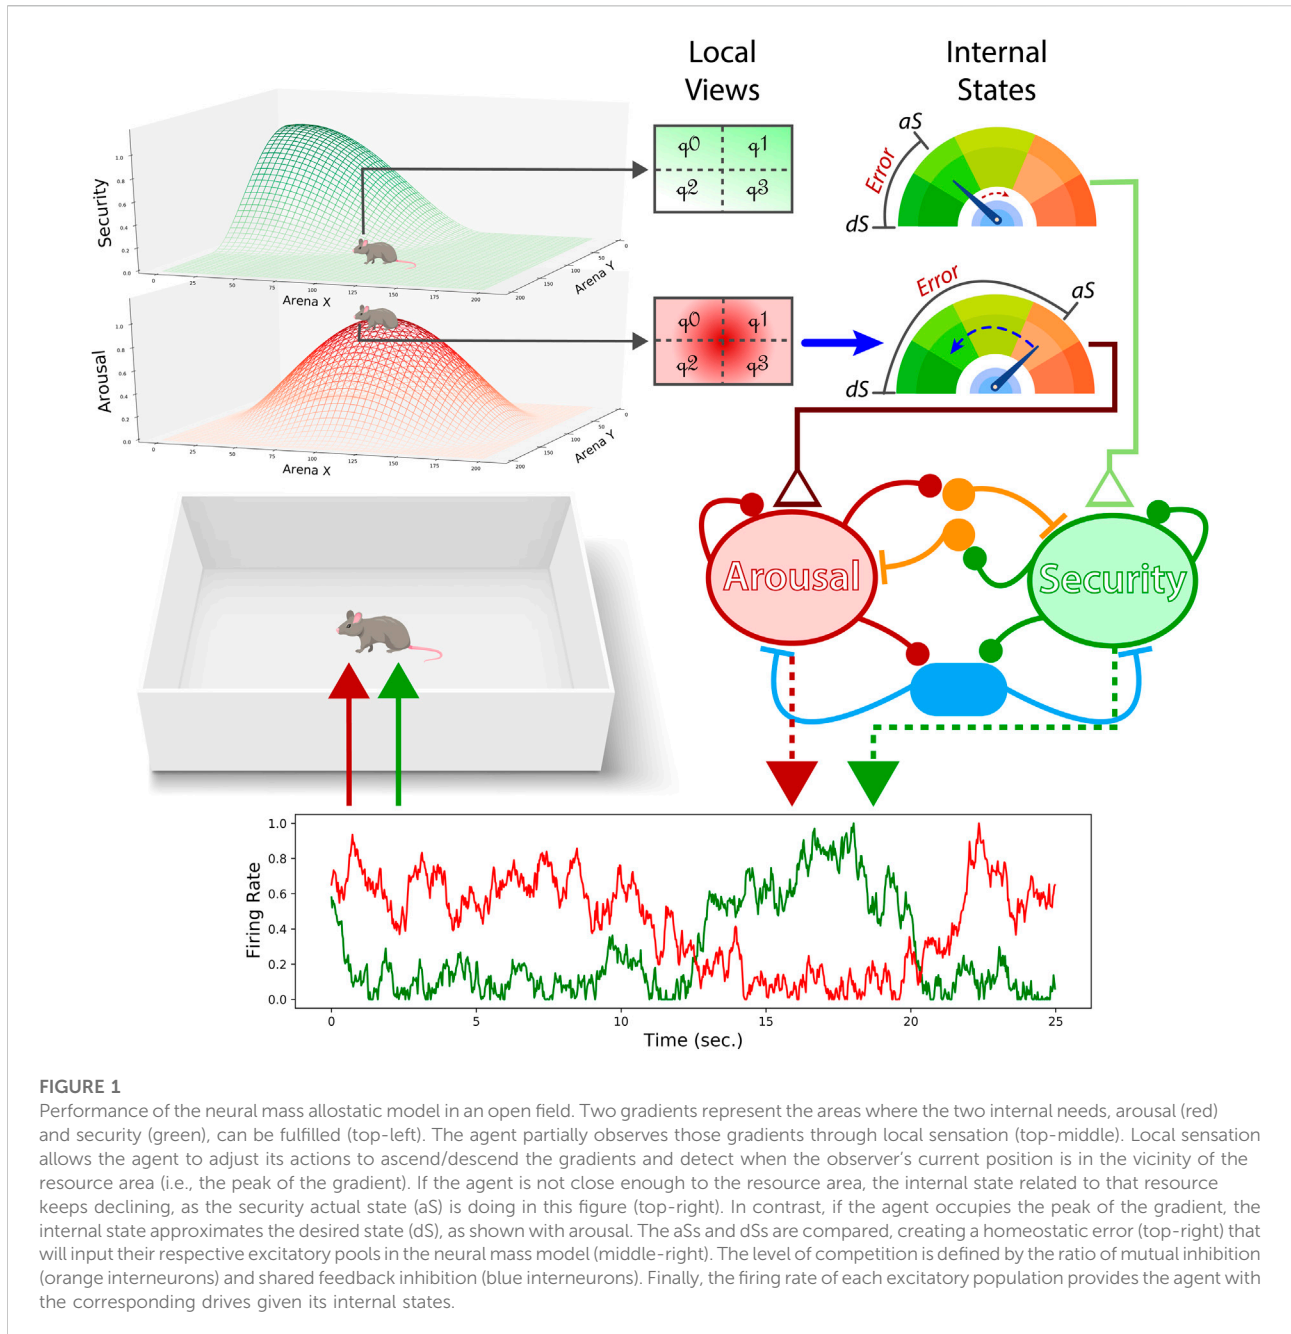
In 2010, Sanchez-Fibla et al. developed what, to our knowledge, is the first computational model of allostatic orchestration (Sanchez-Fibla et al., 2010). The model emulated rodent behavior and physiological states in an open field test (Gould et al., 2009) with simulated and physical robots. This model proposed that the animal's behavior resulted from the interaction between two internal needs: Security, which is fulfilled in one of the arena's corners representing the rodent's home base, and arousal, which would be higher in the center of the arena given the maximum exposure of the animal at that location. However, although reproducing the animal's overall trajectory pattern and occupancy preferences, the model did not elaborate on the neuroscience supporting allostasis.

A more recent model bases the optimal selection of regulatory behaviors on maximizing a subsequent reward (Laurençon et al., 2021). This deployment of a reward-based allostatic model builds on the premises of homeostatic reinforcement learning (HRL) (Keramati and Gutkin, 2014). HRL represents a major refinement of traditional reinforcement learning theories grounding learning protocols on the individual's internal state. HRL successfully explains effects in animal behavior such as alliesthesia, namely, the fluctuations in reward value during resource acquisition (Cabanac, 1971). Still, although HRL leverages temporal and spatial information to improve the self-regulatory strategy, this approach makes a complementary learning process critical for solving the allostatic orchestration problem.

Finally, a third approach suggests that allostatic orchestration emerges from motivational conflict solved *via* attractor dynamics (Jimenez-Rodriguez et al., 2020). In this approach, the attractor dynamics implement competition through cross-inhibition (Usher and McClelland, 2001; Marshall et al., 2015). This framework supports the idea that attractor dynamics can underlie the optimal selection of regulatory behavior and explained their duration and latency. However, an explanation of how neural correlates of internal needs' implement such competing dynamics is unclear.

Grounding the design principles of a model of allostasis in state-of-the-art neuroscience constitutes a pending task for previous modeling approaches that have largely relied on algorithmic solutions. We suggest that novel approaches can overcome this challenge by focusing on the core behavior systems of the mammalian brain (Merker, 2013; Verschure, 2016). Interoception of physiological needs such as hydration, nutrition, thermoregulation, or sleep is generally attributed to distinct specialized hypothalamic nuclei (Strecker et al., 2002; Blouet and Schwartz, 2010; Nakamura, 2011; Zimmerman et al., 2017). In contrast, other psychological needs (e.g., social interaction) depend on more distributed brain networks (Lee et al., 2021). Importantly, recent studies suggest that these nuclei, and so the internal needs they represent, are not independent of each other but hold a competing relationship through inhibitory interactions (Burnett et al., 2016; Osterhout et al., 2022; Qian et al., 2022). This competition between internal drives could serve as the basis of allostatic orchestration by imposing a winner-take-all mechanism represented as attractor dynamics. Thus, irrelevant drives are suppressed, and the singleness principle of action is supported (Sherrington, 1906). This research literature suggests that an attractor-based approach is suitable for modeling allostasis.

In contexts where animals constantly self-regulate multiple internal needs, decision-making could be hampered by attractor forces sustained after drive-completion behaviors. Indeed, *in vivo* studies suggest that cortical areas involved in decision-making operate in a critical regime (close to a phase transition) (Ma et al.,



2019) that occasionally evolves to supercriticality (saturated population response), a regime that supports effective information transmission (Li et al., 2019). Nonetheless, how the system returns to criticality from a supercritical regime is poorly understood. The paraventricular hypothalamic nucleus (PVH) is a good candidate to inhibit the main interoceptive nuclei, recover their basal population activity after drive-completion behaviors, and create the initial conditions for the next cycle of allostatic orchestration. From one side, PVH mediates many diverse motivational functions, including thirst (Zimmerman et al.,

2017), hunger (Blouet and Schwartz, 2010), and thermoregulation (Nakamura, 2011). Conversely, corticotrophin-releasing hormone (CRH) neurons in the PVH are suggested to be sensitive toward reward acquisition. Specifically, PVH CRH neurons get inhibited during drive-completion behaviors, representing a potential source of global inhibition to the rest of hypothalamic interoceptive nuclei (Yuan et al., 2019). In this research work, we model this form of decision reset by applying general inhibitory inputs to the excitatory populations once an internal need has been satisfied.

In the following sections, we will describe a novel neural mass model of allostatic orchestration grounded on the interoceptive mechanisms of the mammalian brain. As in (Sanchez-Fibla et al., 2010), the model will be embedded in an agent endowed with a so-called core behavior system (CBS) (Merker, 2013; Verschure, 2016). In other words, besides sensing its internal environment, a competence attributed to the hypothalamus (Strecker et al., 2002; Blouet and Schwartz, 2010; Nakamura, 2011; Zimmerman et al., 2017), the agent can orientate in an external environment and perform basic navigation based on an appetitive-aversive axis, cognitive functions attributed to the superior colliculus and the zona incerta-periaqueductal gray axis, respectively. The resultant model is tested in both static and dynamic environments. In the static condition, we aim to elucidate if our model can faithfully replicate both previous models of allostatic control and rodent behavior in an open field test (Gould et al., 2009). In the dynamic environment, we will further explore the robustness of our model when environmental opportunities to satisfy internal needs decrease over time. Finally, we will analyze how inducing subcritical dynamics after drive fulfillment facilitates switches in self-regulatory strategies.

Materials and methods

To better understand the potential of competing dynamics between internal drives in facilitating need orchestration and stability of the self, we built a novel allostatic model grounding its design on contemporary research literature. Consequently, a neural mass model was built incorporating two distinct populations sensitive to homeostatic markers while holding a competing relationship (Figure 1). Aiming for convergent validation, we equipped a synthetic agent with this biologically-constrained model and analyzed its ability to defend its internal states by navigating an external environment.

Homeostatic systems

Assuming that living organisms have interoceptive capabilities to assess their internal state, we conceptualized each internal need as a homeostatic mechanism. Here, actual and desired states are compared providing a measure of homeostatic error (that is, $hE_i = |dS_i - aS_i|$). Importantly, homeostatic systems have no power to provide regulatory responses by themselves. Instead, they bias those behaviors by constantly feeding the neural mass allostatic model with homeostatic errors (Figure 1).

The transition between actual internal states responded to the following dynamical law,

$$aS_i(t) = aS_i(t-1) - dR_i(t) + rI_i(t) \quad (1)$$

where $aS_i(t-1)$ is the actual state of the homeostatic system i in the immediate previous timestep, $dR_i(t)$ corresponds to the decay rate applied to each homeostatic system at time t , and $rI_i(t)$ represents the resource impact at time t . Importantly, resource impact will be 0 unless the agent is in the vicinity of the resource. By applying different decay factors and resource impacts, internal states can evolve with different temporal dynamics.

This dynamical law, similarly applied in (Jimenez-Rodriguez et al., 2020), notably differs from the methods applied in (Sanchez-Fibla et al., 2010), where the agent drew its internal states directly from gradients mapped onto the two-dimensional arena.

The neural mass allostatic model

Wilson-Cowan equations modified as in (Amil and Verschure, 2021) were used to model two need-sensitive neural populations (Figure 1). This modification allowed to account for mutual and shared feedback inhibition held between the excitatory populations, as follows:

$$\begin{aligned} \tau \frac{dD_1(t)}{dt} = & -D_1 \\ & + f(w_+D_1 + hE_1 - Qw_-D_2 - (1-Q)w_=f(D_1 + D_2)) \\ & + \sigma\xi(t) \end{aligned} \quad (2)$$

$$\begin{aligned} \tau \frac{dD_2(t)}{dt} = & -D_2 \\ & + f(w_+D_2 + hE_2 - Qw_-D_1 - (1-Q)w_=f(D_1 + D_2)) \\ & + \sigma\xi(t) \end{aligned} \quad (3)$$

where $f(x)$ is the logistic $f - hE_i$ function,

$$f(x) = \frac{F_{max}}{1 + e^{-\frac{(x-\theta)}{k}}} \quad (4)$$

In these equations, τ is the time constant determining the timescale of population dynamics. D_i is the drive magnitude represented as the mean firing rate of the excitatory population i . w_+ , w_- , and $w_=$ are the weights for recurrent connections within the excitatory population, mutual inhibition, and feedback inhibition, respectively. Q represents the mutual/feedback inhibition ratio, a variable that allows for inducing a controlled level of competition. σ and ξ are the variance and magnitude of Gaussian noise provided to the excitatory populations (See Supplementary Figures S1–S4 for Q and noise parameter search). Finally, F_{max} , k , and θ are the maximum firing rate, gain, and threshold parameters of the $f - hE_i$ logistic curve respectively (Eq. 4).

Orientation

Methods enabling the orientation and navigation of the synthetic agent in the arena are based on (Sanchez-Fibla et al., 2010). In both conditions (static and dynamic environments), we served from three-dimensional gradients to represent the location of resources fulfilling specific internal needs. Notably, gradients are used solely to support navigation, and internal states are not directly linked to the agent's location in the arena, as in the original study. Instead, internal states follow their dynamics, as explained above. The decision to use gradient-based methods for navigation is supported by representations of future navigational goals in the orbitofrontal cortex (Basu et al., 2021).

We implemented a partially observable environment by providing the agent with local sensations of the gradient areas surrounding its position. The local sensation was divided into four quadrants (Figure 1) to allow goal-directed-like navigation. By observing differences between the upper horizontal quadrants (q_i^0 and q_i^1), H_i^{sign} controlled the agent's orientation. AD_i^{sign} implements an appetitive-aversive behavioral axis (gradient ascent/descent) by comparing agent (aGL) and resource gradient location (rGL); understanding aGL as the mean gradient value between the four quadrants and rGL as the gradient value at its peak.

$$H_i^{sign} = \begin{cases} 1 & \text{if } q_i^0 < q_i^1 - th \\ -1 & \text{if } q_i^0 > q_i^1 + th \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

$$AD_i^{sign} = \begin{cases} 1 & \text{if } aGL_i^0 < rGL_i^1 - th \\ -1 & \text{if } aGL_i^0 > rGL_i^1 + th \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

Internally-driven navigation

Outputs of the excitatory populations of the neural mass allostatic model are integrated with H_i^{sign} and AD_i^{sign} orientational signals to result in internally-driven goal-oriented-like navigation, following:

$$Navigation = c + \left(\sum_{i=1}^{nGrad} H_i^{sign} \cdot AD_i^{sign} \cdot Exc_i^{output} \right) \cdot \frac{1}{nNeeds} \quad (7)$$

where c is a constant ensuring a default action going forward, Exc_i^{output} accounts for the output of the excitatory population i of the neural mass allostatic model, and $\frac{1}{nNeeds}$ is a normalization factor given the number of implemented needs.

Experimental design

A first study was conducted to evaluate the competence of the neural mass allostatic model in replicating rodent behavior

during an open field test. In this study, a synthetic agent navigated a static simulated environment to defend its internal states of security and arousal. 50 experiments were carried out to analyze navigational patterns consistency and the internal dynamics of the agent along the simulation.

In the second study, the synthetic agent endowed with the neural mass allostatic model incorporated two distinct internal needs: thermoregulation and hydration. A dynamic simulated environment allowed us to assess the agent's ability to adapt its navigation according to environmental opportunities. 50 experiments were carried out to analyze the consistency of the navigational patterns, the internal dynamics taking place along the simulation, and the relationship between these internal dynamics and environmental changes.

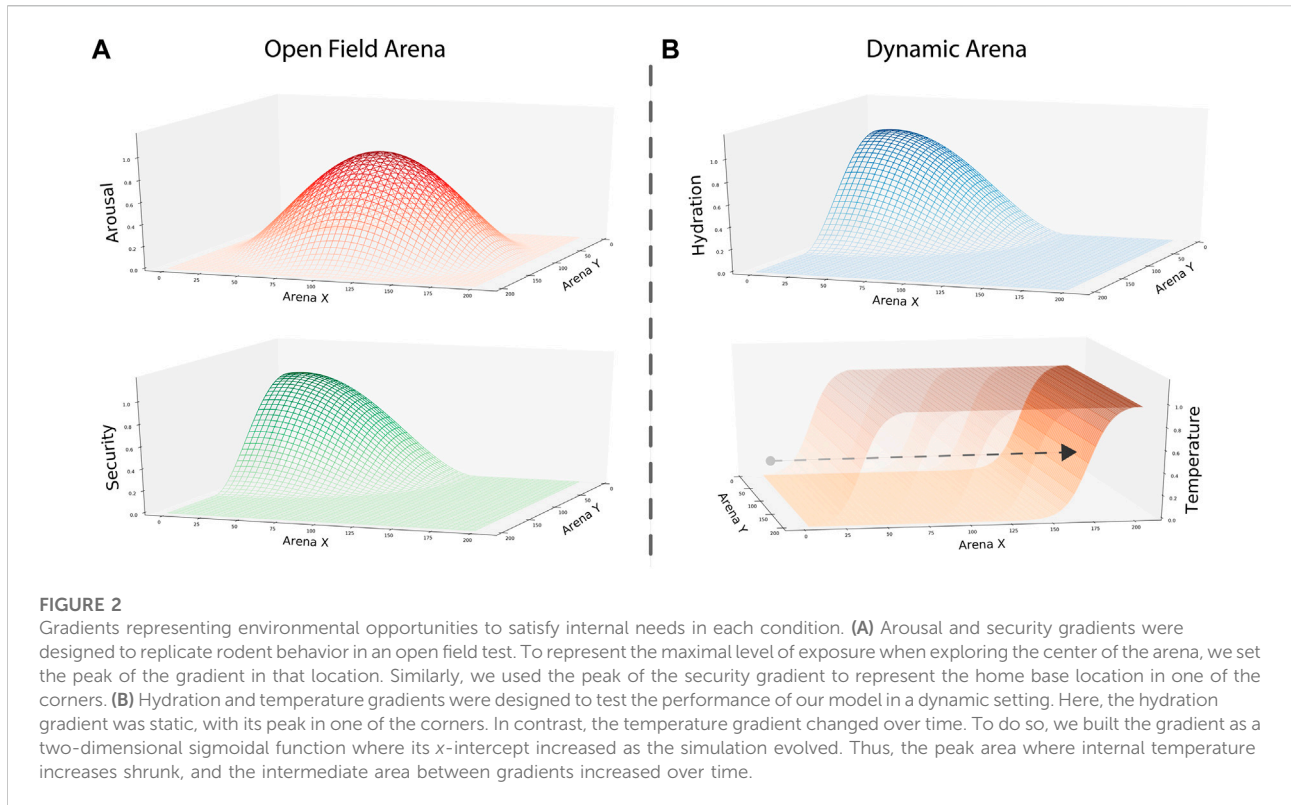
In both the first and the second studies, the agent must constantly decide what internal drives should base its navigation on to maximize stability. This continuous decision-making condition represents a major difference from the original work in which the model's design is based on (Amil and Verschure, 2021) and is a novel challenge to overcome. Therefore, a third study was conducted to evaluate the advantages of inducing subcritical dynamics after need resolution. As in the second study, a synthetic agent navigated a dynamic environment to defend its internal states of hydration and thermoregulation. 50 experiments were conducted to assess the advantages of applying global inhibition after drive-completion behavior compared with the second study, where we did not consider such an effect.

Arenas

The arenas simulated in static and dynamic conditions were designed as two-dimensional 200×200 matrices incorporating one gradient per internal need. These gradients represented the location in the two-dimensional space where each internal need can be satisfied.

In the first study, we built a static environment following the gradient design of the open field test used in (Sanchez-Fibla et al., 2010). Thus, we implemented two gradients mapping the opportunities to calm security and arousal drives. The security gradient was designed as a Gaussian gradient with its peak in the top-left corner, representing the home base location. Meanwhile, the arousal gradient was designed with its peak in the center of the arena, representing the animal's exposure level. The configuration of these gradients aimed to replicate rodent navigational patterns, understood as a preference to occupy the home base and to explore the arena close to the walls (i.e., thigmotaxis) with occasional transversals (Figure 2A).

In the second and third studies, a dynamic environment implemented two gradients to map the opportunities to satisfy thermoregulation and hydration internal needs. We used a similar static gradient to security in the first study to



represent the thirst-calming resource—nonetheless, the gradient representing the area where internal temperature increases changed over time. Designed as a two-dimensional sigmoidal function, the area providing temperature was large at the beginning of the simulation. However, as the simulation evolved, the x-intercept of our sigmoidal gradient increased, shortening the temperature gradient's peak area and establishing a larger gap between the two resource locations (Figure 2B). Thus, at the experiment's end, the temperature gradient's peak area only covers a short area of the arena at its bottom limit.

Efficiency, fairness and stability

To further understand the relationship between internal and external variables, in the second and third studies, we compute three measures traditionally used in game theory (Binmore, 2005; Hawkins and Goldstone, 2016; Freire et al., 2020). Efficiency provides a measure of how good the system is in maintaining the internal states close to their setpoints. Fairness designates whether there is any bias in the engagement of drives. Lastly, stability, computed as the mean square homeostatic error, comprises both the magnitude and the difference between homeostatic errors to provide a general measure of the agent performance.

To inform about the evolution of these measures throughout the experiment, they were calculated for every 1/100 fraction of the duration of the experiment. Hence, mean internal and desired states only represented those elements in the considered fraction across the 50 experiments.

$$Efficiency = \frac{meanT + meanH}{nNeeds} \quad (8)$$

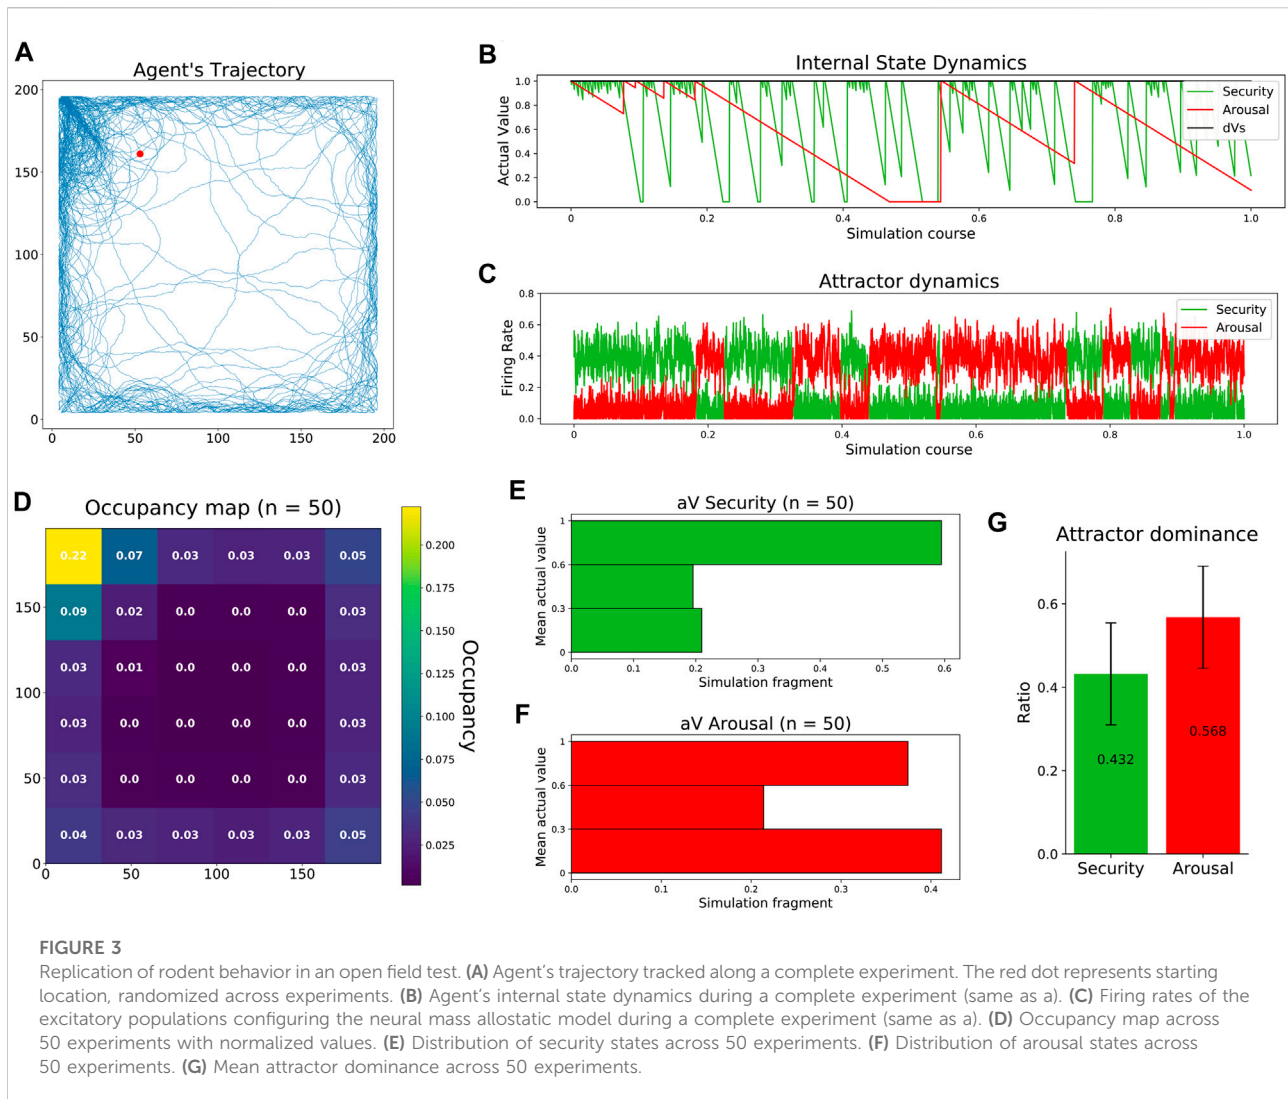
$$Fairness = \frac{|meanT - meanH|}{nNeeds} \quad (9)$$

$$Stability = 1 - \sum_{i=1}^{nGrad} (meanIS_i - meanDS_i)^2 \quad (10)$$

Results

Open field test

In the first study, as previous works did (Sanchez-Fibla et al., 2010), we aim to reproduce the navigational patterns of a rodent in an open field test. As in rodents, the trajectories of our agent showed preferences toward the walls (thigmotaxis) and the top-left corner (home base). At the same time, occasional transversals explored the center of the arena (Figure 3A, Supplementary Figure S5). The observed trajectory patterns resulted from the competition between two drives with different temporal

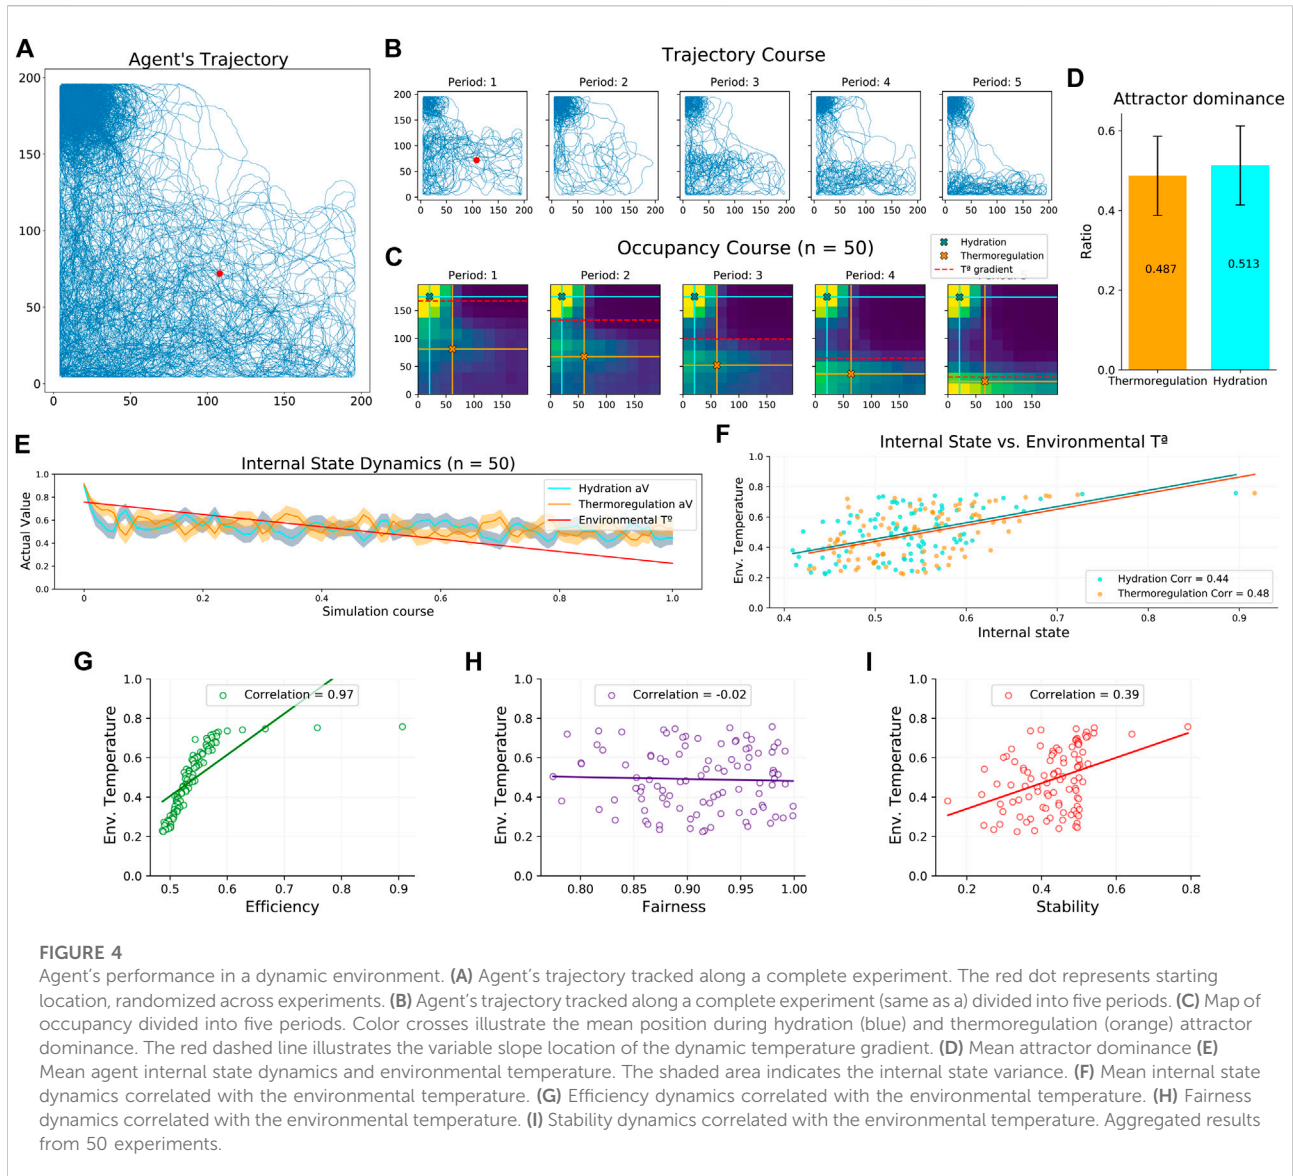


dynamics. Fast-decaying security made the agent constantly revisit its home base, while slow-decaying arousal allowed the agent to occasionally visit the center of the arena (Figure 3B). Competing dynamics between the two internal needs emerged when feeding the excitatory populations of our neural mass model with the corresponding homeostatic errors, allowing the dominant attractor to inhibit its opponent, thus resolving drive orchestration (Figure 3C). After carrying out 50 experiments, agent trajectories showed consistency in their occupancy pattern: a clear preference to visit the home base, navigate close to the walls, and avoid the center of the arena (Figure 3D). The distribution of the internal states during those 50 simulations was very informative. On the one hand, the agent maintained a high level of security during a large part of the simulations (Figure 3E). On the other hand, the state of arousal followed a bimodal distribution indicating the agent was either high or low aroused during the experiment. This bimodal

distribution can be interpreted as follows: First, while the agent pursues security, an aversion toward the center of the arena induces low states of arousal. Then, low internal states, in turn, promote transversals that fully replenish the arousal of the agent (Figure 3F). Finally, attractor dominances, i.e., the simulation period where the firing rate of one excitatory neural pool exceeded the firing rate of the other, showed a balanced activation of both attractors with a slight bias toward arousal (Figure 3G).

Dynamic environment

In the second study, we assess the robustness of our novel neural mass allostatic model in a dynamic environment. In this condition, we observe that the agent's trajectory accurately tracks the environmental gradients (Figure 2B) occupying their peak areas



(Figure 4A, Supplementary Figure S6). Specifically, when dividing the complete experiment into five periods of equal length, we can observe that the agent's trajectory constantly adapted to the changes in the temperature gradient (Figure 4B). Thus, when initializing the experiment (period 1), the peak area of the temperature gradient covers a large part of the arena, and the agent navigates more extensively. However, in the last period (period 5), the peak area is reduced to a thin region close to the bottom border, and the agent's navigation adapts to it. This trajectory pattern was consistent across the 50 experiments, as the occupancy maps suggest (Figure 4C). Spearman correlation analysis indicated that the mean Y axis position of the agent during thermoregulation highly correlated with the temperature gradient's slope location ($r(498) = 0.96$, $p < 0.001$), confirming the agent's trajectory adaptation. Once again,

attractor dominance was balanced (Figure 4D), suggesting that although the environment asymmetrically reduces the opportunities to fulfill the agent's internal needs, the neural mass model imposes a well-balanced competition without neglecting any of the drives. Indeed, the mean internal state across the experiment shows that thermoregulation and hydration levels are well maintained without biases (Figure 4E), and both are equally correlated with the environmental temperature (mean value of the temperature gradient) (Figure 4F). To understand in detail the relationship between internal states and environmental dynamics, we studied this relationship in terms of efficiency, fairness, and stability metrics. Correlating these measures with the environmental temperature, we observed that efficiency strongly depends on the opportunities to fulfill internal needs

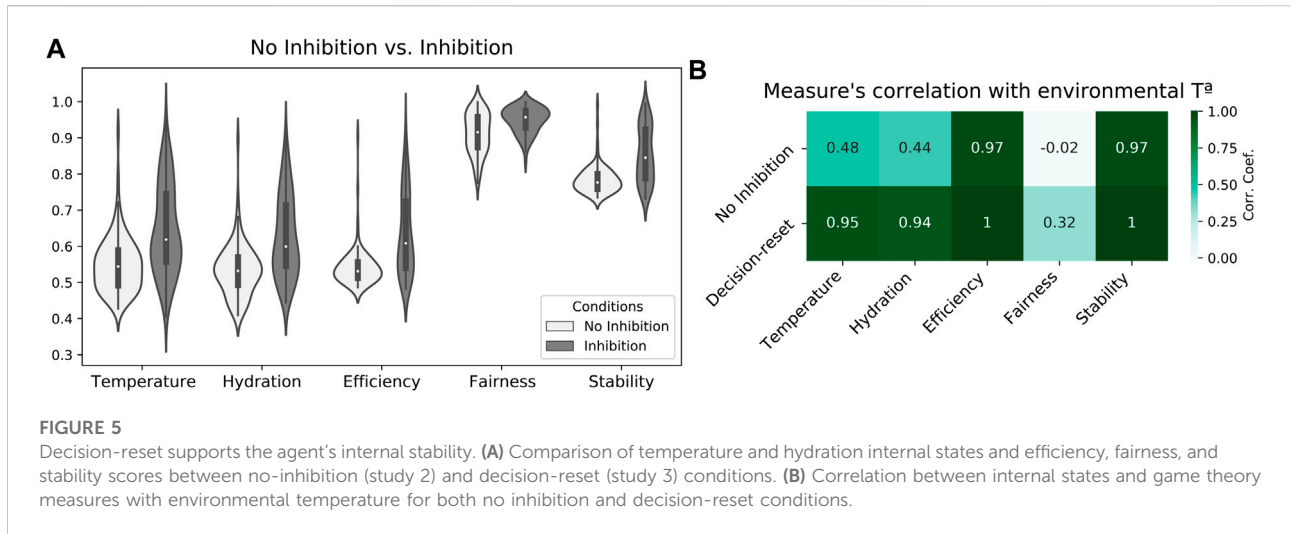


FIGURE 5

Decision-reset supports the agent's internal stability. **(A)** Comparison of temperature and hydration internal states and efficiency, fairness, and stability scores between no-inhibition (study 2) and decision-reset (study 3) conditions. **(B)** Correlation between internal states and game theory measures with environmental temperature for both no inhibition and decision-reset conditions.

(Figure 4G). This result was expected, given the environmental dynamics. When the environmental temperature decreases (i.e., temperature gradient peak shrinks), gradients' peak areas are more distant, forcing the agent to navigate a larger area where its internal states decrease. However, environmental changes do not affect the fairness level (Figure 4H), validating that the balance of attractor dominance is defended even when receiving asymmetric influences from the environment. Lastly, stability reports difficulties minimizing the mean square homeostatic error as temperature decreases (Figure 4I). According to our previous analysis, this result would be better explained by environmental influences on efficiency than fairness.

Criticality-driven decision-reset after decision accomplishment

Finally, a third study explores the potential limitations of attractor-based allostatic models. Specifically, by inhibiting the model's excitatory neural populations (once drive-completion behaviors have been performed), we explore whether sustained attractor forces could hamper individual self-regulation in study 2. Results showed that agents widely benefit from inducing a critical regime to set the initial conditions for each cycle of allostatic orchestration. Specifically, mean internal states along the simulation (calculated every 1/100 fraction of the experiment) across 50 simulations indicated that agents maintained their internal states better when applying decision-reset (Inhibition). (Figure 5A). Consequently, increased internal states also resulted in increased efficiency, fairness, and stability scores. A Mann-Whitney test indicated statistically significant differences when comparing the two conditions in each measure, $U(N_{No-inhibition} = 100, N_{inhibition} = 100), p < 0.001$. Furthermore, this enhanced performance occurred in larger alignment with the

environmental temperature (Figure 5B), which indicates the agent's internal states decreased when the scarcity of resources prevented a better self-regulatory strategy. Altogether, these results suggest that attractor forces sustained after drive-completion behaviors hampered allostatic orchestration, and criticality-driven decision-reset provides an effective mechanism to facilitate a more flexible decision-making process.

Discussion

Previous computational works have contributed to explaining animal self-regulatory behavior (Sanchez-Fibla et al., 2010; Keramati and Gutkin, 2014; Jimenez-Rodriguez et al., 2020). However, these models have not been validated in a broad range of tasks with emphasis on their robustness in the face of varying task conditions, while in parallel, their grounding in the brain mechanisms underlying allostatic orchestration is not fully elucidated. We hypothesize that attractor dynamics originate from the competing relation between different hypothalamic interoceptive nuclei and their loops through the zona incerta, which are further perturbed by the superior colliculus (input and orienting control), and the central gray for triggering reactive species-specific behaviors. To test the adequacy of allostatic attractor dynamics in orchestrating internal needs, we built a biologically-constrained computational model of allostatic orchestration. The resultant model implements competition between internal drives by comprising two drive-sensitive excitatory neural populations that apply mutual and feedback inhibition. To allow a synthetic agent to navigate a virtual external environment based on its internal state, we endowed it with the ability to orientate and perform basic appetitive-aversive navigational behaviors. In this manner, the hypothalamus, superior colliculus, and zona incerta-periaqueductal gray axis are represented in our model, providing a first approximation of the core behavior system (Merker, 2013; Verschure, 2016).

Deploying the allocentric synthetic agent in a virtual environment, we tested the competencies and robustness of the neural mass allostatic model. Specifically, our model allowed the agent to 1) navigate an open field reproducing rodent behavior, 2) adapt its navigation to environmental changes without neglecting any internal needs, and 3) benefit from criticality reset to optimize the interoceptive-driven decision-making process. Altogether, our results supported our hypothesis empirically validating attractor dynamics as an inherent feature of the hypothalamic circuitry that can underlie robust allostatic orchestration. However, the attractor dynamics that emerge in our model need to be further validated by physiological data, ideally by direct neural recordings of the hypothalamic interoceptive nuclei. Lacking ground truth benchmarking data, the temporal dynamics of the homeostatic markers employed in this work were arbitrarily set, representing an additional methodological limitation. Moreover, now that we have assessed the robustness and fidelity of the model to track environmental dynamics, we need to assess its scalability.

Additionally, our work opens new questions in the quest to understand self-regulation. For instance, a better understanding of the interplay between stress and risky internal states is needed. In our studies, the level of competition between neural populations remained fixed at an optimal point. However, the ratio between mutual and shared inhibition could be modulated by stress markers such as acetylcholine (Kawaguchi, 1997). How stress can modulate the competition held between internal drives is a question that remains unanswered. Another open question is how allostatic orchestration interplays with predictive and contextual allostasis to assist each other. We propose to structure allostasis as a cognitive architecture following the four organizational layers of the Distributed Adaptive Control theory of mind (Verschure, 2012). Here, a first somatic level endows the agent with predefined internal needs. Then, the reactive level rapidly orchestrates behaviors directed to calm the most salient drives. In parallel, the adaptive layer learns associations between external and internal events to predict future homeostatic errors and solve them in anticipation. Lastly, a contextual layer enriches the representation of the self-regulatory strategies and enables their memorization.

Finally, contributions to the field of allostasis can expand the boundaries of self-regulation in its understanding and implementation. On the one hand, drives and motivations should also respond to a hierarchical organization that encodes not only priority but also abstraction, virtualization, and replaceability of the drive (Verschure, 2016). Maslow's hierarchy was a first approximation to capture this organization (Maslow, 1943); however, 7 decades later, this theory has not been further advanced in the face of new insights into the dynamics of self-regulation. How different needs within the same priority level organize or how new needs arise and substitute previous ones are questions not answered yet. On the other hand, further investigation could

shed light on how a hierarchical organization of needs fits with a neural mass model. Secondly, general-purpose robot instantiations can benefit greatly from the self-regulatory principles described here. Previous works elaborated on how a multi-agent robotic recycling plant can implement homeostatic and allostatic principles to self-organize at both single-agent and large-scale levels (Rosado and Verschure, 2020). Similar architectures can be designed and implemented to advance a variety of scenarios where robot autonomy is key such as robot delivery or space exploration. Indeed, autonomy in artificial intelligence and robotics places the question of self-regulation at the center of the research of synthetic embodied cognition and consciousness.

Data availability statement

The original contributions presented in the study are included in the article/Supplementary Material, further inquiries can be directed to the corresponding authors.

Author contributions

OGR, AFA, and ITF designed the protocol. OGR conceived and conducted the experiments, analyzed the results, and wrote the manuscript. PV initiated and supervised the research. All authors reviewed and approved the manuscript.

Funding

This study was supported by Counterfactual Assessment and Valuation for Awareness Architecture—CAVAA (European Commission, EIC 101071178) and Hybrid Human-Robot RECYcling plant for electriCal and eLEctRonic equipment—HR-RECYCLER (Horizon 2020, project ID: 820742).

Acknowledgments

We thank our colleagues from the Synthetic, Perceptive, Emotive and Cognitive Systems Lab (SPECS) who provided valuable feedback during the entire conception of the present work. We would also like to show our gratitude to those students from the Cognitive Systems and Interactive Media master (University Pompeu Fabra, Barcelona) who showed interest in the research line and contributed indirectly with parallel research projects. Last but not the least, we are also grateful to the reviewers for taking the time and effort to review the manuscript and provide such valuable comments and suggestions.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated

organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frobt.2022.1052998/full#supplementary-material>

References

- Amil, A. F., and Verschure, P. F. (2021). Supercritical dynamics at the edge-of-chaos underlies optimal decision-making. *J. Phys. Complex.* 2, 045017. doi:10.1088/2632-072x/ac3ad2
- Basu, R., Gebauer, R., Herfurth, T., Kolb, S., Golipour, Z., Tchumatchenko, T., et al. (2021). The orbitofrontal cortex maps future navigational goals. *Nature* 599, 449–452. doi:10.1038/s41586-021-04042-9
- Bernard, C. (1865). *Introduction à l'étude de la médecine expérimentale*. Paris, France: J. B. Baillière et fils.
- Binmore, K. (2005). *Natural justice*. Oxford: Oxford University Press.
- Blouet, C., and Schwartz, G. J. (2010). Hypothalamic nutrient sensing in the control of energy homeostasis. *Behav. Brain Res.* 209, 1–12. doi:10.1016/j.bbr.2009.12.024
- Burnett, C. J., Li, C., Webber, E., Tsaousidou, E., Xue, S. Y., Brüning, J. C., et al. (2016). Hunger-driven motivational state competition. *Neuron* 92, 187–201. doi:10.1016/j.neuron.2016.08.032
- Cabanac, M. (1971). Physiological role of pleasure: a stimulus can feel pleasant or unpleasant depending upon its usefulness as determined by internal signals. *Science* 173, 1103–1107. doi:10.1126/science.173.4002.1103
- Cannon, W. B. (1939). *The wisdom of the body*. New York, NY: W. W. Norton and company INC.
- Freire, I. T., Moulin-Frier, C., Sanchez-Fibla, M., Arsiwalla, X. D., and Verschure, P. F. (2020). Modeling the formation of social conventions from embodied real-time interactions. *PLoS one* 15, e0234434. doi:10.1371/journal.pone.0234434
- Gould, T. D., Dao, D. T., and Kovacsics, C. E. (2009). *The open field test. Mood and anxiety related phenotypes in mice*, 1–20. Berlin, Germany: Springer Science+Business Media.
- Hawkins, R. X., and Goldstone, R. L. (2016). The formation of social conventions in real-time environments. *PLoS one* 11, e0151670. doi:10.1371/journal.pone.0151670
- Jimenez-Rodriguez, A., Prescott, T. J., Schmidt, R., and Wilson, S. (2020). "A framework for resolving motivational conflict via attractor dynamics," in *Conference on biomimetic and biohybrid systems* (Germany: Springer), 192–203.
- Kawaguchi, Y. (1997). Selective cholinergic modulation of cortical gabaergic cell subtypes. *J. neurophysiology* 78, 1743–1747. doi:10.1152/jn.1997.78.3.1743
- Keramati, M., and Gutkin, B. (2014). Homeostatic reinforcement learning for integrating reward collection and physiological stability. *Elife* 3. doi:10.7554/elife.04811
- Laurençon, H., Ségerie, C.-R., Lussange, J., and Gutkin, B. S. (2021). Continuous homeostatic reinforcement learning for self-regulated autonomous agents. *arXiv preprint arXiv:2109.06580*.
- Lee, C. R., Chen, A., and Tye, K. M. (2021). The neural circuitry of social homeostasis: Consequences of acute versus chronic social isolation. *Cell* 184, 1500–1516. doi:10.1016/j.cell.2021.02.028
- Li, M., Han, Y., Aburn, M. J., Breakspear, M., Poldrack, R. A., Shine, J. M., et al. (2019). Transitions in information processing dynamics at the whole-brain network level are driven by alterations in neural gain. *PLoS Comput. Biol.* 15, e1006957. doi:10.1371/journal.pcbi.1006957
- Ma, Z., Turrigiano, G. G., Wessel, R., and Hengen, K. B. (2019). Cortical circuit dynamics are homeostatically tuned to criticality *in vivo*. *Neuron* 104, 655–664.e4. doi:10.1016/j.neuron.2019.08.031
- Marshall, J. A., Favreau-Peigné, A., Fromhage, L., Mcnamara, J. M., Meah, L. F., and Houston, A. I. (2015). Cross inhibition improves activity selection when switching incurs time costs. *Curr. Zool.* 61, 242–250. doi:10.1093/czoolo/61.2.242
- Maslow, A. (1943). *A theory of human motivation*. psychological review no 50. Washington, DC: American Psychological Association.
- Maturana, H. R., and Varela, F. J. (1991). *Autopoiesis and cognition: The realization of the living*, 42. Germany: Springer Science & Business Media.
- Merker, B. (2013). The efference cascade, consciousness, and its self: Naturalizing the first person pivot of action control. *Front. Psychol.* 4, 501. doi:10.3389/fpsyg.2013.00501
- Nakamura, K. (2011). Central circuitries for body temperature regulation and fever. *Am. J. Physiology-Regulatory, Integr. Comp. Physiology* 301, R1207–R1228. doi:10.1152/ajpregu.00109.2011
- Osterhout, J. A., Kapoor, V., Eichhorn, S. W., Vaughn, E., Moore, J. D., Liu, D., et al. (2022). A preoptic neuronal population controls fever and appetite during sickness. *Nature* 1–8, 937–944. doi:10.1038/s41586-022-04793-z
- Papies, E. K., and Aarts, H. (2016). Automatic self-regulation: From habit to goal pursuit. *Handbook of self regulation: Research, theory, and applications* (New York: Guilford Press).
- Pavlov, I. P. (1955). *Selected works*, 1. Russia: Foreign Languages Publishing House.
- Qian, S., Yan, S., Pang, R., Zhang, J., Liu, K., Shi, Z., et al. (2022). A temperature-regulated circuit for feeding behavior. *Nat. Commun.* 13, 4229–4317. doi:10.1038/s41467-022-31917-w
- Rosado, O. G., and Verschure, P. F. (2020). "Distributed adaptive control: An ideal cognitive architecture candidate for managing a robotic recycling plant," in *Conference on biomimetic and biohybrid systems* (Germany: Springer), 153–164.
- Sanchez-Fibla, M., Bernardet, U., Wasserman, E., Pelc, T., Mintz, M., Jackson, J. C., et al. (2010). Allostatic control for robot behavior regulation: a comparative rodent-robot study. *Adv. Complex Syst.* 13, 377–403. doi:10.1142/s0219525910002621
- Sherrington, C. S. (1906). *The integrative action of the nervous system*, 35. New Haven: Yale University Press.
- Sterling, P. (2012). Allostasis: a model of predictive regulation. *Physiology Behav.* 106, 5–15. doi:10.1016/j.physbeh.2011.06.004
- Sterling, P. (1988). "Allostasis: a new paradigm to explain arousal pathology," in *Handbook of life stress, cognition and health*. Hoboken, New Jersey: John Wiley and Sons.

- Sterling, P. (2020). *What is health?: Allostasis and the evolution of human design*. Cambridge: MIT Press.
- Strecker, R., Nalwalk, J., Dauphin, L., Thakkar, M., Chen, Y., Ramesh, V., et al. (2002). Extracellular histamine levels in the feline preoptic/anterior hypothalamic area during natural sleep-wakefulness and prolonged wakefulness: an *in vivo* microdialysis study. *Neuroscience* 113, 663–670. doi:10.1016/s0306-4522(02)00158-6
- Tschantz, A., Barca, L., Maisto, D., Buckley, C. L., Seth, A. K., and Pezzulo, G. (2022). Simulating homeostatic, allostatic and goal-directed forms of interoceptive control using active inference. *Biol. Psychol.* 169, 108266. doi:10.1016/j.biopsycho.2022.108266
- Usher, M., and McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* 108, 550–592. doi:10.1037/0033-295x.108.3.550
- Verschure, P. F. (2012). Distributed adaptive control: a theory of the mind, brain, body nexus. *Biol. Inspired Cogn. Archit.* 1, 55–72. doi:10.1016/j.bica.2012.04.005
- Verschure, P. F. (2016). Synthetic consciousness: the distributed adaptive control perspective. *Phil. Trans. R. Soc. B* 371, 20150448. doi:10.1098/rstb.2015.0448
- Wiener, N. (1948). *Cybernetics*. New York: John Wiley & Sons.
- Yuan, Y., Wu, W., Chen, M., Cai, F., Fan, C., Shen, W., et al. (2019). Reward inhibits paraventricular crh neurons to relieve stress. *Curr. Biol.* 29, 1243–1251.e4. doi:10.1016/j.cub.2019.02.048
- Zimmerman, C. A., Leib, D. E., and Knight, Z. A. (2017). Neural circuits underlying thirst and fluid homeostasis. *Nat. Rev. Neurosci.* 18, 459–469. doi:10.1038/nrn.2017.71