

Complementary interactions between classical and top-down driven inhibitory mechanisms of attention

S.C. Low^a, V. Vouloutsi^a, P.F.M.J. Verschure^{a,b,*}

^a *Laboratory of Synthetic, Perceptive, Emotive and Cognitive Systems (SPECS), Institute for Bioengineering of Catalonia (IBEC), The Barcelona Institute of Science and Technology, Barcelona, Spain*

^b *Institució Catalana de Recerca i Estudis Avançats (ICREA), Barcelona, Spain*

Received 18 November 2020; accepted 6 December 2020

Available online 18 December 2020

Abstract

Selective attention informs decision-making by biasing perceptual processing towards task-relevant stimuli. In experimental and computational literature, this is most often implemented through top-down excitation of selected stimuli. However, physiological and anatomical evidence shows that in certain situations, top-down signals could instead be inhibitory. In this study, we investigated how such an inhibitory mechanism of top-down attention compares with an excitatory one. We did so in a neurorobotics context where the agent was controlled using an established hierarchical architecture. We augmented the architecture with an attentional system that implemented top-down attention biasing as connection gains. We tested four models of top-down attention on the simulated agent performing a foraging task: without top-down biasing, with only excitatory top-down gain, with only inhibitory top-down gain, and with both excitatory and inhibitory top-down gain. We manipulated the reward-distractor ratio that was presented and assessed the agent's performance using accumulated rewards and the latency of the selection. Using these measures, we provide evidence that excitatory and inhibitory mechanisms of attention complement each other.

© 2020 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Keywords: Selective attention; Inhibition; Foraging; Embodied cognition

1. Introduction

For an agent acting in all but the most simplistic environments, it likely receives more sensory information than it has the capacity to process in the moment. Such prioritisation of task-relevant stimuli over others is called selective attention. It is, however, arguably not an end goal itself, but rather a means to several ends (Allport, 2016) as it is ultimately behaviour that leads to survival or success in

achieving the agent's goals. As behaviours are elicited through action selection based on perceived stimuli, biasing sensory processing would lead to changes in behavioural outcome as well. Selective attention thus has an impact on embodied agents' behaviour, and it is therefore a common method to test the performance of attention models by implementing them in artificial agents.

There are two advantages to implementing models of attention in artificial agents. Firstly, and many a times this alone justifies the effort, an agent equipped with some form of selective attention is better able to interact with more complex environments and exploit greater sensor resolution and sensitivity. This is clearly beneficial in applications

* Corresponding author at: UPC Campus Diagonal-Besòs Building C, Floor 6 Av. d'Eduard Maristany, 10, 14 Barcelona 08930 Spain.

E-mail address: pverschure@ibecbarcelona.eu (P.F.M.J. Verschure).

like robotics and computer vision. Secondly, taking a more ontological perspective, it considers the model not simply as an algorithm solving an abstract problem but something that enables a body to function, which is a key principle of embodied cognition (Anderson, 2003; Engel et al., 2013; Wilson, 2002). Thus, there are a considerable number of implementations of selective attention in artificial agents or robots. An influential, saliency-based model of bottom-up attention was proposed by Itti et al. (Itti et al., 1998) and subsequently inspired many others (Driscoll et al., 1998; Frintrop et al., 2003; Ouerhani and Hügli, 2005). There have also been proposals that include top-down elements of attention to facilitate more sophisticated behaviours, for example through drive management (Breazeal and Scassellati, 1999), cuing from human partners (Haasch et al., 2005), and exploiting a collection of historical bottom-up input to guide top-down attention (Colombini et al., 2016).

The success of these implementations of selective attention supports not only its functional value but also the physiological, experimental literature that has grown alongside this field (e.g. Cutrone et al., 2014; Desimone, 1998; Sani et al., 2017). Nonetheless, common task paradigms tend to revolve around task-relevant stimuli being more easily anticipated than task-irrelevant ones. This has potentially led to a bias towards excitatory top-down models that target task-relevant stimuli. For example, only 5 of the 52 reviewed cognitive architectures included the suppression of task-irrelevant stimuli, space or features; all of them exhibited at least one element of selection and restriction (Kotseruba and Tsotsos, 2020), which is excitatory in nature.

However, there are situations where an excitatory mechanism is inefficient, such as when task-relevant stimuli are difficult to predict but task-irrelevant stimuli are not. Here, inhibition is more suitable as a mechanism of attention, as it targets task-irrelevant stimuli. Indeed, there is strong support for both mechanisms in the mammalian brain (Phillips et al., 2016; Zikopoulos and Barbas, 2012; Wimmer et al., 2015) and they appear to have strengths and weaknesses that are complementary rather than competitive. We therefore investigated how these two mechanisms, both separately and combined, modulated the behaviour of an artificial, autonomous agent performing a controlled, foraging task.

The agent was placed in an environment populated with two types of objects—rewards and distractors—distinguishable by their colours. By approaching and touching an object, the agent effectively demonstrated that its control dynamics had selected that object over other objects in the vicinity. In this study, the agent’s control architecture was based on the biologically grounded Distributed Adaptive Control (DAC; (Pfeifer and Verschure, 1992; Verschure, 2016)), which has been shown to, among other cognitive functions, robustly learn sequences of state-action couplets that eventually lead to a goal state (Marcos et al., 2014). We demonstrated the utility of top-

down attentional biasing in the task, and additionally benchmarked the inhibitory model against the excitatory one. We showed how, between the excitatory and inhibitory models, excitatory biasing led to quicker trials but lower yield of rewards. Most interestingly, we observed that the two mechanisms working together resulted in the best performance across conditions.

2. Methods

2.1. Agent and environment

We placed an embodied agent (Fig. 1) in a virtual arena to perform a foraging task. The environment was created using Python (Rossum, 1995); PyGame (Shinners, 2011) and Box2D (Catto, 2011). The simulated agent was a mobile, two-wheeled robot with seven sensors for proximity and vision each. Both types of sensors projected a ray which returned either the distance or the colour of the first object it hit (Fig. 1-A).

The objects were solid-coloured circles. The distractors were intentionally valued to be neutral instead of punishments (as punishments necessitate aversive behaviour and are task-relevant); a distraction itself has no benefit or loss and is therefore task-irrelevant. The top-down attentional mechanisms tested here were the excitatory model, conceptually congruent with most enhancement-based models of attention, and the inhibitory model, where activity representing task-irrelevant stimuli is suppressed. In both cases, top-down biasing was dependent on the motor decision from the contextual layer (see Control Architecture below) in the previous time step. If the selected action was towards a reward, the excitatory model would excite subsequent signals from the associated sensor while the inhibitory model




Fig. 1. Experimental setup. (A) Example of how the agent senses objects in its environment. The agent has seven visual sensors, represented by the blue circles spread across its front. The white lines show the linear projections of three visual sensors, and the individual sensor is filled with the colour of the object that it detects. A sensor filled with black indicates that it has not detected any object. (B) The first experiment consisted of two stimuli presented to the agent, and the order of the reward (green) and distractor (red) was randomised. The positions of the objects were symmetric about the agent’s midline, but their distance to the midline, indicated by d , was not fixed (0.23 m – 0.78 m). (C) The second experiment consisted of seven objects presented to the agent, and the proportion of rewards was varied as an experimental condition (reward-distractor ratio). The positions at which the objects were placed were fixed to be evenly distributed, but the order was randomised. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

would inhibit it in the next time step(s). This led to a biasing of subsequent sensory competition and, by extension, perceptual information that is next fed into the contextual layer for action selection. The strength of the top-down excitatory and top-down inhibitory signals were defined by their gains, G_{TDexc} and G_{TDinh} respectively (Fig. 2).

2.2. Task description

The agent started every trial with a number of objects equidistant to it and within its field of vision. Regardless of experimental condition, there was always bottom-up sensory competition in the form of intralayer inhibition (W_{intra}). However, the values of the top-down biasing gains (G_{TDexc} and G_{TDinh}) were changed between conditions.

Two experiments, variations of the same task, were run; in both, a trial ended when the agent touched an object and success was determined as the selection of a reward. For the first experiment, the agent was presented with two objects, a reward and a distractor, located equidistant to it but with varying distances to each other (Fig. 1-B). Green indicated a reward and red a distractor. The exact colour of the objects was randomly drawn from a range of intensities and the objects' order was randomised. To determine the effect of each mechanism on its own, one of the top-down biasing gains was kept constant at 0 while

the other was tested in the range [0–0.7] with step sizes of 0.1. 50 trials were run per gain condition.

In the second experiment, the agent was presented with seven objects located equidistant to it. The reward-distractor ratio was varied (Fig. 1-C). Although the positions were fixed, the objects' order was randomised every trial. The objects were evenly distributed in front of the agent over 140° . When active, the top-down biasing gains for both excitatory and inhibitory mechanisms were set to 0.5 in this experiment as, in the first experiment, it had yielded approximately the mean performance for nonzero top-down gains. There were four attentional conditions: with no top-down biasing ($g_{G,TDexc} = g_{R,TDinh} = 0$), with only inhibitory top-down biasing ($g_{G,TDexc} = 0$, $g_{R,TDinh} = 0.5$), with only excitatory top-down biasing ($g_{G,TDexc} = 0.5$, $g_{R,TDinh} = 0$), and with both excitatory and inhibitory biasing ($g_{G,TDexc} = g_{R,TDinh} = 0.5$). In addition, there were four reward-distractor ratios: with one, three, four, and six rewards. Hence, there were 16 conditions, and for each condition 50 trials were run leading to a total of 800 trials.

2.3. Control architecture

As a cognitive architecture, the Distributed Adaptive Control theory (DAC) posits that cognition arises from the interaction between interconnected control loops operating at four increasing levels of abstraction for hierarchical control: soma, reactive, adaptive, and contextual. At the lowest level of abstraction, the soma layer defines the body of the agent in its environment. The layer above it is the reactive control of behaviour, which associates certain sensory input with hard-wired behaviours. Although limited in functionality, the reactive layer is still a necessary part of the architecture (Nolfi, 2002).

Behaviour generated by the reactive layer bootstraps learning of the agent's environment and the sensory outcomes of its actions, which takes place in the adaptive layer. This helps the agent develop behavioural responses that are more noise resistant. In the adaptive layer, these pairs of perceived states and motor action are associated together as state-action couplets, which are then sent to the contextual layer for higher-level cognitive functions like goal selection and planning. Through Bayesian decision-making (Bayes, 1763; Beck, 2008), DAC's contextual layer captures a knowledge level description of intelligence and the principle of rationality by exploiting perceptual and behavioural learning through interaction with the environment (Verschure et al., 2003). This is supported by a long-term memory that is modelled after what is found in animals, which was proposed to accumulate knowledge and abilities to optimise subsequent planning and behaviour (Newell, 1994). In this study, it was manifested by the agent comparing its current state with those in its long-term memory to identify potential goal states, anticipate the state-action couplets that would bring it there, and trigger goal-oriented behaviours.

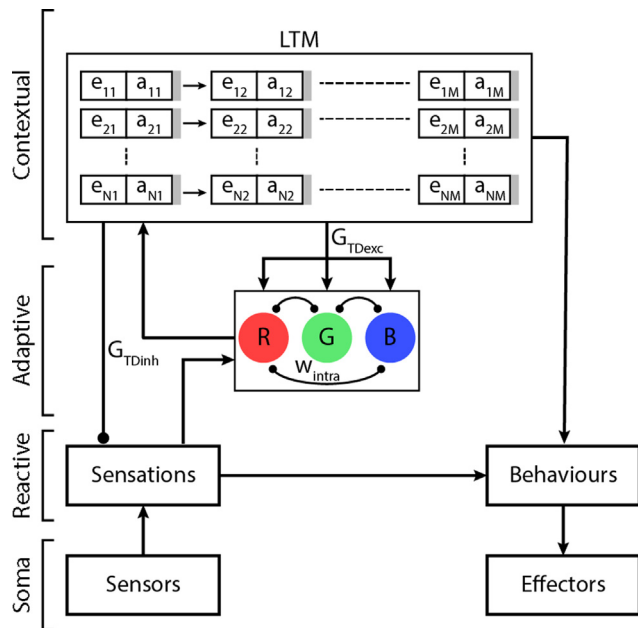


Fig. 2. Control architecture of the agent. Lines terminating with arrowheads show excitatory connections while lines terminating with circles show inhibitory connections. The contextual layer decided on actions and change the gains of the top-down attentional biasing based on its long-term memory (LTM), which is shown here as a collection of sequences of state-action couplets. The visual channels are indicated by the nodes R, G, and B, and the connections which were modulated in the experiments are labelled G_{TDexc} and G_{TDinh} . The gain of the inhibition, representing bottom-up processes, is shown by W_{intra} and applied to all the inhibitory connections between the visual channels.

The agent's behaviour was a combination of motor commands from the reactive and contextual layers (Fig. 2). Although the DAC architecture allows for motor control via the adaptive layer, in this implementation it did not provide any direct motor control to avoid confounding the results. The agent's reactive layer was hard-coded to avoid colliding with walls using the activation of proximity sensors, similar to Braitenberg vehicles (Braitenberg, 1986). The agent also had a reactive drive to randomly explore the environment when it did not detect any objects, ensuring that the agent proactively interacted with the environment.

The adaptive layer integrated bottom-up and top-down signals to produce the perceptual state that was used by the contextual layer. The perceptual information was driven by sensory input, which consisted of three colour channels per visual sensor. The competition observed between sensory cortical neuronal populations (Reynolds et al., 1999) was simulated here with the sensors' activity contributing to the activity of each channel's node. Their activity was weighted by W_{intra} . After this interaction, the winning channel was decomposed into the different contributing sensors, and a winner-takes-all mechanism selected the most active sensor's activity as the perceptual input to the contextual layer.

In conditions with top-down biasing, the top-down gains (G_{TDexc} and G_{TDinh}) were nonzero. The contextual layer matched the perceptual input to the states in the long-term memory. In this way, it predicted the potential reward of following a sequence of motor actions and made decisions that maximised future rewards. This top-down biasing signal is applied to the sensory input of the adaptive layer in the next time step. The contextual layer was pre-populated with appetitive behaviour towards both red and green objects, though their values were not identical, and the same long-term memory was used in all experiments. Although all sequences in the long-term memory terminated with a goal state, they were not necessarily of

the same length. The agent's reactive exploratory behaviour was overridden by the contextual layer, but the reactive collision avoidance took priority over the contextual layer.

3. Results

3.1. Top-down attentional biasing increased selectivity of rewards

In the first experiment, chance level reward selection was 0.5. Both mechanisms improved performance to at least 0.7 (Fig. 3-A). The trials were then separated into three conditions and their durations, in time steps, were analysed: without top-down biasing (68 ± 32 steps), with nonzero inhibitory gains (80 ± 38 steps), and with nonzero excitatory gains (65 ± 33 steps). Trials with nonzero excitatory weights were significantly faster than the other conditions (Table 1).

3.2. Attentional mechanisms led to quicker trials when there were more rewards than distractors

When there was one reward, the top-down attentional mechanism did not significantly affect trial duration (Fig. 3-B; $42. \pm 6.0$ steps (no TD), $41. \pm 7.8$ steps (inh), $40. \pm 8.2$ (exc), $46. \pm 9.3$ steps (both); Friedman, $p = 0.077$, $\chi^2(3) = 6.9$). However, when there were three rewards, we found that the duration of trials with only excitatory top-down biasing was significantly quicker than the other conditions (Friedman, $p = 0.0021$, $\chi^2(3) = 14.$; Table 1). With four rewards, all the conditions with any attentional biasing had similar durations and were all significantly faster than when there was no attentional biasing ($49. \pm 2.1$ steps (no TD), $47. \pm 1.8$ steps (inh), $46. \pm 2.3$ steps (exc), $46. \pm 2.9$ steps (both); Friedman, $p = 0.0039$, $\chi^2(3) = 13.$). When there were six rewards, there were sig-

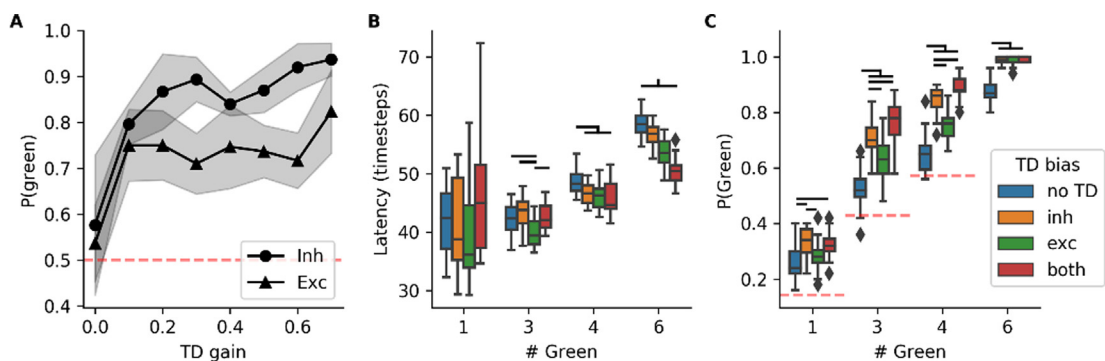


Fig. 3. Performance of the agent in the two experiments. (A) In the first experiment, the probability of a successful trial, determined by the agent approaching and selecting a reward (green) when the top-down biasing mechanism was either only inhibitory or only excitatory. The red line indicates chance level. (B) The mean durations per trial in each condition in the second experiment. Within each column, the top-down biasing groups are ordered, from left to right, as follows: no top-down biasing (no TD), only inhibitory top-down biasing (inh), only excitatory top-down biasing (exc), and both inhibitory and excitatory top-down biasing (both). (C) The probability of the agent selecting a reward in the second experiment. The order and labelling are the same as in the previous plot. The red dotted lines indicate chance level for that reward proportion. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 1

Post-hoc Monte Carlo permutation testing statistics for mean differences in latencies of trials. Exp: experiment; N(G): number of green objects, which serve as rewards; CI (95%): the 95% confidence interval.

Exp	N(G)	Control	Test	Difference	CI (95%)		p-value
					lower	upper	
1	1	No TD	Exc	-3.1	-6.0	-0.25	0.043
	1	No TD	Inh	13.	9.8	16.	<0.001
	1	Exc	Inh	16.	14.	18.	<0.001
2	3	No TD	Inh	1.2	-0.44	2.9	0.15
	3	No TD	Exc	-2.2	-3.6	-0.57	0.0096
	3	No TD	Both	0.55	-0.93	2.2	0.50
	3	Exc	Both	2.7	1.3	4.1	0.0012
	3	Inh	Exc	-3.4	-4.8	-1.8	<0.001
	4	No TD	Inh	-2.0	-3.2	-0.88	0.0024
	4	No TD	Exc	-2.4	-3.8	-1.1	0.0018
	4	No TD	Both	-3.0	-4.5	-1.5	<0.001
	4	Inh	Exc	-0.42	-1.7	0.81	0.53
	6	No TD	Inh	-1.8	-3.0	-0.71	0.006
	6	No TD	Exc	-4.9	-6.3	-3.6	<0.001
	6	No TD	Both	-8.2	-9.4	-6.9	<0.001
	6	Exc	Both	-3.3	-4.6	-1.8	<0.001
	6	Inh	Exc	-3.1	-4.4	-1.8	<0.001

nificant differences between all the attentional conditions ($59. \pm 2.0$ steps (no TD), $57. \pm 1.9$ steps (inh), $54. \pm 2.4$ (exc), $50. \pm 2.2$ steps (both)); Friedman, $p < 0.001$, $\chi^2(3) = 47.$).

3.3. Combining excitatory and inhibitory attentional mechanisms improved performance regardless of reward proportion

When there was one reward, the excitation model did not perform significantly different from the model without

top-down biasing (Fig. 3-C, Table 2). The models with only inhibitory biasing and both excitatory and inhibitory biasing performed similarly, and at a level significantly better than without top-down biasing and with only excitatory biasing (0.33 ± 0.054 (inh), 0.32 ± 0.049 (both); Friedman, $p = 0.001$, $\chi^2(3) = 16.$; Table 2)

The models with some form of top-down attention all performed better than the one without any top-down biasing once the reward-distractor ratio increased. The inhibitory only model performed better than the excitatory only model, and performance was best when both inhibi-

Table 2

Post-hoc Monte Carlo permutation testing statistics for mean differences in probability of success. Exp: experiment; N(G): number of green objects, which serve as rewards; CI (95%): the 95% confidence interval.

Exp	N(G)	Control	Test	Difference	CI (95%)		p-value
					lower	upper	
2	1	No TD	Inh	0.076	0.040	0.11	<0.001
	1	No TD	Exc	0.033	-0.0030	0.066	0.071
	1	no TD	Both	0.073	0.040	0.10	<0.001
	1	Inh	Exc	-0.043	-0.074	-0.008	0.015
	1	Inh	Both	-0.003	-0.033	0.028	0.85
	3	No TD	Inh	0.19	0.14	0.23	<0.001
	3	No TD	Exc	0.10	0.053	0.15	<0.001
	3	no TD	Both	0.25	0.20	0.29	<0.001
	3	Inh	Exc	-0.082	-0.13	-0.039	0.001
	3	Inh	Both	0.064	0.020	0.10	0.006
	4	No TD	Inh	0.19	0.15	0.23	<0.001
	4	No TD	Exc	0.11	0.062	0.14	<0.001
	4	no TD	Both	0.23	0.19	0.27	<0.001
	4	Inh	Exc	-0.088	-0.12	-0.053	<0.001
	4	Inh	Both	0.039	0.013	0.068	0.013
	6	No TD	Inh	0.12	0.097	0.13	<0.001
	6	No TD	Exc	0.11	0.094	0.13	<0.001
	6	no TD	Both	0.12	0.10	0.14	<0.001
	6	Inh	Exc	-0.002	-0.013	0.006	0.63
	6	Inh	Both	0.004	-0.003	0.010	0.18

tory and excitatory mechanisms were used. With three rewards, this effect was significant (0.52 ± 0.078 (no TD), 0.70 ± 0.061 (inh), 0.62 ± 0.084 (exc), 0.77 ± 0.076 (both); Friedman, $p < 0.001$, $\chi^2(3) = 43.$) and with four rewards as well (0.65 ± 0.073 (no TD), 0.84 ± 0.051 (inh), 0.76 ± 0.056 (exc), 0.88 ± 0.040 (both); Friedman, $p < 0.001$, $\chi^2(3) = 43.$). When there was a large majority of rewards, there was a ceiling effect but the model without top-down biasing remained significantly poorer in performance (0.87 ± 0.042 (no TD), 0.99 ± 0.012 (inh), 0.99 ± 0.018 (exc), 0.99 ± 0.0098 (both); Friedman, $p < 0.001$, $\chi^2(3) = 45.$; Table 2).

4. Conclusion

For agents acting in dynamic environments, it is important to prioritise task-relevant stimuli for processing. Selective attention facilitates this, and typically it is thought to do so via top-down excitation of neurons representing selected stimuli or features while top-down inhibition of close competitors sharpens the contrast in neural activity between selected stimuli and the rest (Itti and Koch, 2001). However, there is increasing interest in an attentional mechanism that instead selects stimuli and features for inhibition. We investigated the effectiveness of the two mechanisms of attention, using an artificial agent that performed a simple foraging task. In this task, the environment consisted of rewards and distractors and the goal of the agent was to reach a reward. The evaluation of the different attentional models was based on the agent's performance (proportion of successful trials) and the time needed to reach an object. Furthermore, the models were evaluated with a variety of reward-to-distractor ratios.

We found that all combinations of attentional biasing led to faster trials with increasing proportions of rewards to distractors. At first glance, it might seem that top-down attentional processes prevent the agent from being distracted by other rewards. However, the contextual layer was the main driver of behaviour even in conditions without top-down attention. As the contextual layer exhibited decision inertia, where previously-selected sequences were prioritised over others, all decisions made would exhibit decision inertia regardless of the top-down biasing condition. Decision inertia alone is thus insufficient to explain the reduced trial durations in conditions with top-down biasing, suggesting that top-down attentional mechanisms do more than simply ensuring that a decision is carried out till its goal is met.

The agent performed better with top-down attentional biasing than without for all reward-to-distractor ratios, demonstrating the usefulness of top-down attention even in a task as simple as this. Most notably, the attentional model that consistently performed the best was one that combined both excitatory and inhibitory mechanisms. This supports the hypothesis that the two mechanisms are complementary.

While the current setup was sufficient for our purposes, it still is a form of visual search of task-relevant stimuli. Future work could test the models in a paradigm that is closer to an oddball task, where what is predictable is also task-irrelevant. In addition, the visual sensors used in the agent were very low resolution. While they approximate vision in humans, it is nowhere close to the complexity of visual input from human eyes. Its simplicity, and that of the environment and objects, allowed the sensory competition to be simulated cleanly through representations of three colour channels. As future work, it would be interesting to test the models' performance in more ecological environments and with a wider range of stimulus features as perceptual states. Additionally, the calculation for bottom-up saliency can be more detailed in terms of biological-grounding, for example including a parameter for contrast. Also, learning and attention are closely coupled (e.g. Eldar et al., 2013; Grossberg, 1999), and a potentially interesting field of research would be how the excitatory and inhibitory models of attention could lead to differences in the learning process either through changes in neuronal synapses or through higher-level learning strategies.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This project was funded by H2020 Research and Innovation program (#787061, ANITA), ERC H2020 (#840052, Cognitive RGS), H2020-Research and Innovation action EU.3.1.5.3 (#826421, VBC), and H2020-EU.2.1.5.1. (#820742, HR-Recycler). There are no conflicts of interest to declare.

References

- Allport, A. (2016). Selection for action: Some behavioral and neurophysiological considerations of attention and action. *Perspectives on Perception and Action*, 395–419.
- Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial Intelligence*, 149(1), 91–130.
- Bayes, T. (1763). LII. An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, FRS communicated by Mr. Price, in a letter to John Canton, AMFR S. *Philosophical Transactions of the Royal Society B*(53), 370–418.
- Beck, J. M. et al. (2008). Probabilistic population codes for Bayesian decision making. *Neuron*.
- Braitenberg, V. (1986). Vehicles: Experiments in synthetic psychology. *Philosophy Review*.
- Breazeal, C., & Scassellati, B. (1999). A context-dependent attention system for a social robot. *IJCAI international joint conference on artificial intelligence*.
- Catto, E. (2011). Box2d: A 2d physics engine for games.
- Colombini, E. L., da Silva Simões, A., & Ribeiro, C. H. C. (2016). An attentional model for autonomous mobile robots. *IEEE Systems Journal*, 11(3), 1308–1319.

- Cutrone, E. K., Heeger, D. J., & Carrasco, M. (2014). Attention enhances contrast appearance via increased input baseline of neural responses. *Journal of Vision*, 14(14), 16.
- Desimone, R. (1998). Visual attention mediated by biased competition in extrastriate visual cortex. *Philosophical Transactions of the Royal Society*, 1245–1255.
- Driscoll, J. A., Peters, R. A., & Cave, K. R. (1998). Visual attention network for a humanoid robot. In *IEEE international conference on intelligent robots and systems* (pp. 1968–1974).
- Eldar, E., Cohen, J. D., & Niv, Y. (2013). The effects of neural gain on attention and learning. *Nature Neuroscience*.
- Engel, A. K., Maye, A., Kurthen, M., & König, P. (2013). Where's the action? The pragmatic turn in cognitive science. *Trends in Cognitive Sciences*.
- Frintrop, S., Rome, E., Nüchter, A., & Surmann, H. (2003). An attentive, multi-modal laser 'eye'. *Lecture notes in computer science (including subseries lecture notes in artificial intelligence and lecture notes in bioinformatics)*.
- Grossberg, S. (1999). The link between brain learning, attention, and consciousness. *Consciousness and Cognition*.
- Haasch, A., Hofemann, N., Fritsch, J., & Sagerer, G. (2005). A multi-modal object attention system for a mobile robot. *2005 IEEE/RSJ international conference on intelligent robots and systems*.
- Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*.
- Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11), 1254–1259.
- Kotseruba, I., & Tsotsos, J. K. (2020). 40 years of cognitive architectures: Core cognitive abilities and practical applications. *Artificial Intelligence Review*.
- Marcos, E., Ringwald, M., Duff, A., Sánchez-Fibla, M., & Verschure, P. F. M. J. (2014). The hierarchical accumulation of knowledge in the distributed adaptive control architecture. *Computational and Robotic Models of the Hierarchical Organization of Behavior*.
- Newell, Allen (1994). *Unified theories of cognition*. Harvard University Press.
- Nolfi, S. (2002). Power and limits of reactive agents. *Neurocomputing*.
- Ouerhani, N., & Hügli, H. (2005). Robot self-localization using visual attention. *Proceedings of IEEE international symposium on computational intelligence in robotics and automation*.
- Pfeifer, R., & Verschure, P. (1992). Distributed adaptive control: A paradigm for designing autonomous agents. In *Toward a practice of autonomous systems: proceedings of the first European conference on artificial life* (pp. 21–30). Cambridge, MA, USA: MIT Press.
- Phillips, J. M., Kambi, N. A., & Saalman, Y. B. (2016). A subcortical pathway for rapid, goal-driven, attentional filtering. *Trends in Neurosciences*, 39(2), 49–51.
- Reynolds, J. H., Chelazzi, L., & Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience*, 19(5), 1736–1753.
- Rossum, G. (1995). Python reference manual.
- Sani, I., Santandrea, E., Morrone, M. C., & Chelazzi, L. (2017). Temporally evolving gain mechanisms of attention in macaque area V4. *Journal of Neurophysiology*, 118(2), 964–985.
- Shinners, P. (2011). PyGame.
- Verschure, P. F. M. J. (2016). Synthetic consciousness: The distributed adaptive control perspective. *Philosophical Transactions of the Royal Society B*.
- Verschure, P. F. M. J., Voegtlin, T., & Douglas, R. J. (2003). Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature*, 425(6958).
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4), 625–636.
- Wimmer, R. D., Schmitt, L. I., Davidson, T. J., Nakajima, M., Deisseroth, K., & Halassa, M. M. (2015). Thalamic control of sensory selection in divided attention. *Nature*, 526(7575), 705–709.
- Zikopoulos, B., & Barbas, H. (2012). Pathways for emotions and attention converge on the thalamic reticular nucleus in primates. *Journal of Neuroscience*, 32(15), 5338–5350.