

Research Article

Open Access

Hanno Hildmann*, Miguel Almeida, Abdel F. Isakovic, and Fabrice Saffre

Focusing attention in populations of semi-autonomously operating sensing nodes

<https://doi.org/10.1515/pjbr-2019-0008>

Received November 14, 2018; accepted January 15, 2019

Abstract: Cognition and the cognitive processing of sensory information in biological entities is known to occur over multiple layers of processing. In the example of human vision there are a vast number of photo-receptors feeding into various layers of cells which pre-process the original information before it arrives to the brain (as biased data). We propose to use a mechanism known to theoretical biologists as a means to bring about adaptive self-organization in colonies of social insects, and to apply it to such early stage signal processing. The underlying mathematical model is simple, and in the coming years, robotics will move into an era when aggregating simple computation devices into massively large collectives becomes feasible, making it possible to actually build such distributed cognitive sensing systems.

Keywords: artificial perception, nature-inspired optimisation, autonomous robotics, artificial intelligence, cell level cognition, swarm intelligence, cognitive psychology, theoretical biology

1 Introduction

The ability to control focus and granularity of sensory input is arguably at the foundation of human cognition, and, possibly less controversially, of critical importance for the

advancing field of autonomous robotics. Social cognitive theory distinguishes among three modes of agency [1]: direct (personal), proxy (relying on the actions of others) and collective (requires *socially coordinative and interdependent effort*). The same can be said about entities of lower cognitive ability (e.g., cells) or collections of cyber-physical systems (e.g., sensors). There, *direct* is the unit's own data generation, *proxy* the unit's ability to focus due to its neighbours providing additional data to cover for otherwise neglected areas and *collective* refers to the performance of the population of units as a whole.

In biological systems large numbers of cells commonly operate together in a way that every unit acts independently but in reaction to - and interaction with - its neighbours. When cells act as sensing nodes in sensory organs, this results in multiple layers of information processing where a large number of information is filtered and processed before it is finally communicated to the higher level information processing units (i.e., the brain). This efficient (albeit biased) multi-layer signal pre-processing may become necessary for complex cyber-physical systems. For example, in a smart city, increasingly large numbers of connected sensing devices allow the central control system to provide situational awareness to civil defense and public safety forces.

However, the sensing capabilities of any equipment, biological or mechanical, is bounded. Increasing the level of detail (e.g., the resolution of a camera or the focus of a collection of visual receptors) means reducing the area that can be covered.

Allocating increased attention to multiple areas or objects in a field of perception *almost always* [2] results in reduced data quality somewhere else (simply due to the overall capability of the sensing device being bounded). The concentration of sensing ability (at the cost of reduced perception elsewhere) is known as *focusing* [2]. Exactly how this works is not entirely understood yet [3]. What is clear, though, is that there are some processes happening at a local level where changes are initiated through the interaction of neighbouring sensing nodes, while others are subject to outside cues originating from more complex analysis of the perceived data.

***Corresponding Author: Hanno Hildmann:** TNO, Oude Waalsdorperweg 63, 2597 AK Den Haag, The Netherlands; NEC Research Labs Europe (NLE), Kurfürsten-Anlage 36, 69115 Heidelberg, Germany; BrainCreators, Prinsengracht 697, 1017 JV Amsterdam, The Netherlands; E-mail: hanno@cypherpunx.org

Miguel Almeida: NEC Research Labs Europe (NLE), Kurfürsten-Anlage 36, 69115 Heidelberg, Germany

Abdel F. Isakovic: Cornell University, Ithaca, NY 14850, USA; Khalifa University of Science and Technology, P.O. box 127788 Abu Dhabi - United Arab Emirates

Fabrice Saffre: Khalifa University of Science and Technology, P.O. box 127788 Abu Dhabi - United Arab Emirates; BT Applied Research, Adastral Park, Martlesham Heath, IP5 3RE, United Kingdom

We propose a simple approach to carry out signal pre-processing at a local level. This paper is inspired by cognitive psychology [4]. The processing of visual information in the human brain, e.g., happens through numerous filtering processes [4] before data arrives in the brain. This requires selective processing of *some* information [5]. We know this to be achieved using both, low-level grouping (cf. *Gestalt-Theory*, [4]) as well as high-level cues [3] (such as e.g., expecting to see two eyes in a face, symmetrically located above and around a nose) which may require a level of conscious interpretation of the perceived image.

In cognitive psychology, this is referred to as the *problem of visual attention* [2], [5]. Humans [6] and primates [7] have been shown to have the ability to focus attention of individual cells consciously, but this ability is limited ($\approx 6\%$ in case of some monkeys [6]); most of the attention focusing is happening at a less-than conscious level.

Our approach is based on models from theoretical biology (for social insects such as ants and termites [8]). The mechanism was successfully applied to small collectives of UAVs for bio-security applications.

We previously reported on related work, which is currently subject to a (pending) patent application [9], in [10].

2 Background & previous work

2.1 Scope and context

Theories and mathematical models on how attention is focused in living animals have been proposed by experts in the field (e.g., [2], [3]) and it is not the aim or claim of this paper to compete with those. The mechanism we propose relies on simple mathematical operations to drive self-organization in an artificial sensing system. In this system, the individual sensing nodes interact only with their neighbours yet collectively ensuring that an entire spectrum or area is covered, while at the same time facilitating the focusing of attention (resolution, measured granularity, data quality, etc) on specific areas.

An important feature of the approach is the underlying simplicity which lends itself for implementation in devices with very low computational power. Furthermore, due to the distributed nature of the approach and the fact that the devices only react to their immediate neighbours (bounded computational cost), the approach is expected to scale well with growing population size.

Motivated by the encouraging performance results from testing the UAV swarm, we present our investigations into the performance of the approach when applied to larger collectives of semi-autonomously operating entities.

2.2 Inspiration

Nature-inspired approaches have been successfully applied to labor division [11], resource allocation [12] or scheduling [13]. Perhaps some of the most efficient nature-inspired solutions can be found in the areas of cognition, in general, and focusing, in particular, where our traditional mechanistic, linear approaches fail. In the field of robotics, especially in the context of the design of autonomous robots [14], ideas from biology and self-organization have generated increasing interest.

Our work is inspired by the behavior of social insects, living in colonies of many thousands of individuals. Such colonies of e.g., bees [15], termites [16] or ants [17] can operate as a whole and in semi-stable pattern until some catalyzing event occurs [18], causing a localised behaviour change that addresses the impact of the event while leaving the remainder of the colony operating normally.

Mathematical models for these collective have been successfully applied to e.g., load-balancing [19] and resource-allocation [12]. In line with our previous work, the approach investigated in this article relies on insights gained from the study of termites (cf. [8]).

Without going into too much detail, termites use feedback from the environment to guide their probabilistic behaviour. In simple terms, their decision paradigm could be described as '*rich gets richer*' or '*rich gets poorer*'. In other words, agents either locally amplify or balance aggregations of building materials, purely based on the situation on site. By doing so continuously and many times over, stable constellations are favored and - in an unchanged environment - the collective settles into a workable (i.e., constraint satisfying) pattern. As long as changes in the environment are not occurring too dramatically, the colony can also adapt to these changes and - as a whole - operates within acceptable limits.

2.3 Previous work

In this paper we apply aforementioned decision paradigm to large collectives of (computationally) very simple sensing nodes. This is analogous to previous work where we applied this to far more complex agents with encouraging results. We use the remainder of this page to outline this previous work but refer the interested reader to [10] for technical specifications and implementation details. Without further ado:

2.3.1 Application scenario

A swarm of collaborating agents, each with their own sensing equipment (and the ability to tune this), is tasked with providing situational awareness for an area too large for any individual to cover alone. To bound the communication overhead, communication is only required between agents sharing coverage of a data stream.

NEC Europe has investigated UAV-based animal health control in New Zealand using IR-cameras to provide information about the body temperatures of individual animals in large herds (roaming freely in the wild).

2.3.2 Computational agents

We used custom built drones capable of outdoor flight operations (cf. Figure 1, left). All basic flight operations and navigation dynamics are performed by a Pixhawk flight module (Figure 1, right, bottom), the onboard computer running our algorithm is a Raspberry Pi 2 (Figure 1, right, top). The autopilot software facilitated the simulation of flight operations, making it possible for us to operate additional instances of agents on other Raspberry Pis to form a hybrid swarm of up to 25 drones.

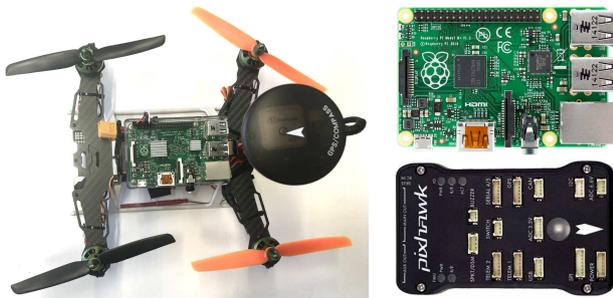


Figure 1: (left) a drone prototype (used in actual flight operations), (right, top) the Raspberry Pi 2 and (right, bottom) the pixhawk flight module / autopilot.

2.3.3 Implementation and evaluation

The drones were built to facilitate real world testing of the algorithm using only the on-board computing hardware. Various technical and societal concerns and challenges (cf. [20]) make it extremely difficult to deploy seizable UAV swarms *in the wild*, which lead us to consider swarms comprised of real as well as simulated devices. All devices operated separately and independently on their respective

physical hardware. This enabled us to perform our proof-of-concept trial and to generate the data presented in [10], which we summarise briefly in Section 4.1.

Due to the practical nature of the project within which the application was built, the members of a swarm, albeit identical, were physically separated and fully independent units, communicating using ROS over VPN. Each software agent was operating individually, blissfully unaware of its embodiment (i.e., a software agent does not *know* whether it is connected to a real drone or a simulation of a drone).

3 Mathematical model and algorithm

We model a collection of simple signal processing agents or *cells*. These can be organised in an n -dimensional plane; in this article we use a 2D plane as this matches the arrangements of physical sensors in the real world but this simplification is inconsequential for the applicability of the approach ($n \geq 2$).

Cells are signal processing devices but this could be of very low complexity: without loss of generality we can use the extremely simple example of a cell reporting only the average of the signal values from its inputs.

Cells have variable fields of *coverage* and can in some rudimentary way affect which signals are allocated to them. Overlapping coverage indicates that the respective cells both have access to the signals received there. Figure 2 illustrates this for 14 cells with somewhat circular fields of coverage.

The complexity of the signal processing performed by the individual cells is bounded only by their processing power. In our model, all signals processed by a cell are processed the same way. Cells operate at a specific level of *attention*, which directly translates to the number of signals a cell covers: because all signals receive equal amounts of processing power, reducing the number of signals increases the amount each signal receives. In this setting, a linear change in the radius of the coverage results in non-linear change of the attention allocated to each signal within that coverage.

One measure of performance is the summed up amount of attention all signals are receiving. Given the mentioned non-linearity, this means that cells should converge to similar attention levels (subject to the topology). Cells with rudimentary autonomous decision making ability might want to temporarily increase attention for some signals (e.g., a spike in a signal might be triggering this re-

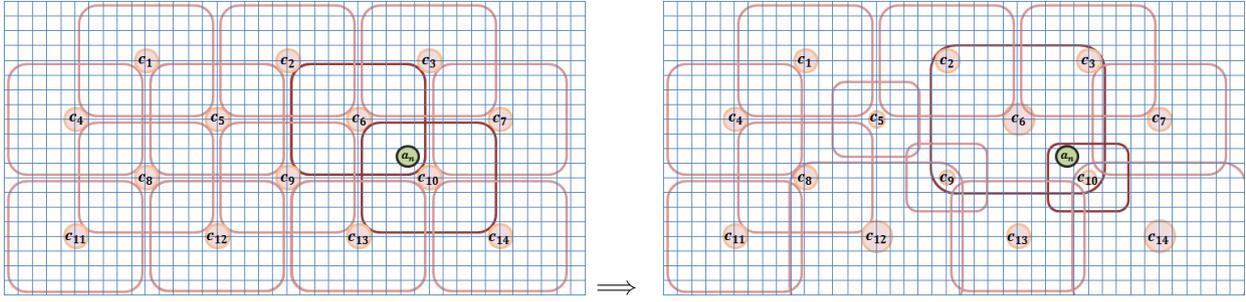


Figure 2: A graphical representation of overlapping areas of coverage of “cells” (processing units). In the above example, cells c_6 and c_{10} can both process signal a_n and with all cells operating at the same *focus* level (left), both of them do. In that case, the decision which cell processes signal a_n is inconsequential and simply subject to agreement between c_6 and c_{10} . If, however, signal a_n requires more attention, then the processing cell has to focus (narrow its coverage), which inevitably results in reducing the number of signals this cell can continue to process (right). To ensure continuous processing of all signals, other cells have to assist by taking on some of the otherwise orphaned signals. This can have a ripple effect, affecting even cells further away. Note: we consider the (more difficult) view where all signals processed by a cell receive equal amounts of processing power but the approach also works if this is not the case.

action). To focus attention, a cell would have to contract its coverage, thereby increasing the quality of the signal processing across all signals in the (reduced) coverage. This can be a temporary change, which is undone after closer investigation of said signal is completed. As long as this change is in effect, we include the deviation from the desired focus into our performance measure.

3.1 Mathematical model

The mechanism requires a population C of sensing nodes, or cells, c ; $C = \{c_1, \dots, c_n\}$. The topology t of these cells is such that each cell c_i is neighbouring a limited number of other cells: $t(c_i) = \{c_j, \dots, c_k\}$, where the term *neighbouring* captures that two cells share access to at least one specific signal a . If the population were viewed as edges in a communication graph, to be neighbours would mean to *be connected*. The fact that cells are only connected locally is important for the scalability of the approach.

This population C , as a whole, processes a set A of signals a ; $A = \{a_1, \dots, a_m\}$. We denote the fact that cell c_i is processing signal a_j by $o(c_i, a_j)$. Note: whether a signal can be processed by more than one cell has implications for the application but is not relevant for our model (see Section 5 for a discussion) as long as there is one dedicated cell *owning* this signal. We take $o(c_i, a_j)$ to express this ownership.

Cells have a threshold for the acceptable level of their average signal processing quality. This value does not have to be constant across C : in [21] we investigate the impact of different thresholds on our model. Whenever a cell’s threshold is not met, the *excitement* value e_c of the cell increases. Furthermore, if additional attention or sensory

resolution is required (due to e.g., an unexpected change in the signal), then this further adds to the excitement value of this cell.

The exchange of ownership for signals between cells is stochastic, i.e., based on weighted probability determined by the cells’ excitement e_c . The probability of re-allocating a_n from c_i to c_j (i.e., changing $o(c_i, a_n)$ to $o(c_j, a_n)$) is:

$$P_{c_i, c_j}(a_n) = \frac{e_{c_i}}{e_{c_i} + e_{c_j}} \quad (1)$$

If c_i is more excited than c_j , the probability of c_i handing over a signal is above 50%, and the larger the difference in excitement, the more likely it is to hand over ownership of a signal. We can therefore use the excitement value of a cell to affect how likely it is to reduce the number of signals it owns. Periodic increase of excitement will force the system to move out of stable constellations (ownership allocations) resulting in naturally occurring ‘gazing’ (wandering of focus in case of the absence of threshold violation).

3.2 Practical considerations

Cells change their focus by increasing or decreasing their coverage. We simplify cell coverage to being circular. For a cell to reduce the radius of its coverage all of the signals at the fringe have to be handed over to another cell first.

Conversely, if a cell increases its radius then it can thereafter accept all signals at the new fringe (meaning that the actual processing of any of these signals is possible because the cell is already operating at the required focus level). In our simulations signal handover is restricted: cells can only (a) hand over their signals at the outer rim of their coverage and (b) accept signals at, or just outside, their outer rim.

If handing over a signal requires a decrease in focus this will result in a changed excitement value for this cell. When calculating the probability of handing a signal to a cell we use the excitement value that this cell would have *if it already had the signal*. Practically speaking, we maintain two excitement values for each cell: the real one, as well as a hypothetical one representing the situation after the exchange. **Note:** this does not affect the overall performance of the approach other than (greatly) speeding it up since it reduces the number of signals to exchange to those signals that can actually be handed over.

3.3 Communication

To avoid signal loss some level of inter-cell communication is necessary. In its most basic form cells need to know (a) their neighbours, (b) their neighbours excitement values as well as (c) which signal can be handed over to which neighbour. The first two can be easily performed by cells regularly broadcasting their excitement value to their immediate neighbours. This does not even require knowledge of who these neighbours are. However, as signals can only be handed to cells that can provide coverage for them, it helps to attempt handovers only with neighbours that can - at least theoretically - accept a signal. In our simulation, this information is hard-coded into the cells. This does not invalidate the applicability of the idea. In biological systems this can be the case and robotics applications may be built to either include this knowledge up front or with the ability to infer it. The latter allows the topology to be dynamic (i.e., subject to change). **Note:** (c) is not strictly required but does speed up the process significantly by avoiding handover attempts guaranteed to fail.

3.4 Algorithm

Equation 1 defines the probability of handing over a signal from one cell to another. The approach can be implemented to either have cells *hand over* signals to one of their neighbours or, conversely, to *steal* a signal from one of their neighbours. Both approaches have the same effect on the system, with Algorithm 1 we provide an implementation for the former.

By keeping the interaction between cells limited to two cells we can reduce the complexity of the decision making process to a boolean decision (see Equation 2). Since cells pick their interaction partners randomly (see Algorithm 1) this implies that many interactions between neighbours

may have to happen before a signal is allocated to the ideal cell.

This effectively constitutes a trade-off between efficient reallocation of signals and the simplicity of the involved math combined with the number of required exchanges. Given that we propose the approach for devices with (extremely) low computational power this does suit our interests. It does, however, mean that the approach is assumed to run continuously (with the `wait` statement in Algorithm 1 determining the frequency through the control parameter `time_constant`). We expect this to be applied to dynamic scenarios where a continuous solution space exploration is desired. Our algorithm and approach addresses just that.

Algorithm 1 Algorithm for cell c

```

while true do
  wait (time_constant +  $\epsilon$ )
  pick random signal  $a_i$ 
  pick random neighbour  $c_{new}$  capable of processing  $a_i$ 
  read excitement  $e_{c_{new}}$  from  $c_{new}$ 
  if  $f(e_{c_{new}}, e_c) == \text{true}$  (see Equation 2) then
     $o(c_{new}, a_i) = \text{true}$ 
     $o(c, a_i) = \text{false}$ 

```

The approach is inherently decentralised, each cell can act independently of the other cells. Cells waiting to become active can passively accept ownership of signals from an active neighbour. Possible conflicts such as being contacted for two handovers at the same time are resolved by simple protocols (such as e.g., voiding both requests). Varying the waited time ever so slightly (`time_constant` + ϵ , see Algorithm 1) prevents cells from ever synchronizing.

Keeping the underlying math as simple as possible is crucial for the proposed idea. A straight forward and computationally cheap way to implement Equation 1 is shown below. Basically, the decision whether to hand over a signal can be reduced to (implemented as) deciding whether a random value is larger than the excitement value of the active cell:

$$f(e_{c_1}, e_{c_2}) = \text{rand}(0, (e_{c_1} + e_{c_2})) > e_{c_2} \quad (2)$$

We can envision even simpler realizations of Equation 1, maybe even some that do not require programmable hardware. Very simple electronic or biological hardware could be designed to realise the idea of a weighted probability tipped in favor of the cell that has a higher excitement value, resulting in a boolean (yes/no) decision.

4 Results and discussion

The approach discussed in this paper is theoretical and we discuss its application to large population of sensing nodes with very limited processing capabilities. While we foresee many application areas for this, we currently do not have such a system ourselves. Due to this, the results presented are all collected in simulation environments. The idea originates in practical and applied work though, and we know the underlying mechanism to work in a variety of applications related to load balancing [19] and dynamic resource-allocation [12]. The mechanism was developed, implemented and evaluated for complex hardware [10] and the first set of results discussed below (Section 4.1) is taken from that project.

What these results did not provide, however, was insights into the performance of the system for large collectives of nodes or for implementations where minimizing the computational load was paramount. To investigate these aspects, a software implementation was developed that simulated large collectives of devices and facilitated collective data collection. Results from this investigation are presented in Section 4.2.

4.1 Practical field test (small population, powerful computers)

The proposed method was used in a real-world project related to using UAVs for collective outdoor sensing tasks. The approach was implemented to run on the on-board computers of custom made UAVs (cf. Section 2.3), which we flight-tested outside. The results discussed here were collected from a swarm in which the devices operated independently from each other. Due to legal and practical restrictions the full swarm of 25 drones was never flown outdoors at the same time. Instead, the swarm was realised in a simulation environment capable of combining real UAVs with simulated drones. This enabled us to form a swarm of 25 devices, with simulated drones operating on a Raspberry Pi but not actually flying.

Our swarm was deployed over an area of 20×20 locations. UAVs were positioned at the intersection of 4 such locations and distributed homogeneously over the area in a square lattice formation. This setup, albeit unrealistic in real world scenarios, meant that in the absence of an special resolution requirements the optimal solution has the swarm hovering at one altitude with as little coverage overlap as possible.

The general ability of the swarm to independently converge towards good solutions was tested by deploying the UAVs at maximum altitude with the expectation that they would all descend to the same altitude where they delivered optimal data feeds. Therefore, instead of considering the specific altitudes of all 25 UAVs individually we can simply focus on the standard deviation of their altitudes: the closer standard deviation is to zero, the closer all drones in the swarm are to flying at the same altitude.

And indeed, the swarm converged into a low and efficient altitude allocation for all UAVs with the standard deviation dropping to zero. This ability to converge towards an efficient (and in our case: the known optimal) allocation suggests that the approach can be used for our first goal, which is to use it to continuously optimise signal allocation to cells.

Our second goal is to enable a population of cells to react to increased stimuli, i.e., to increase the attention a receiving cell can allocate to a specific signal. In the practical field test with our UAVs, this translates to requesting increased resolution for specific areas from the swarms collective video feeds. This would break the coherent altitudes of the members of the swarm as some devices are required to lower their altitude to deliver the requested resolution. As this comes with the loss of coverage, other UAVs would in turn be forced to rise higher so as to mitigate this coverage loss. As described in more detail in [10], this is exactly what we observed. In addition to efficiently responding to the requirements for increased resolution, the swarm's members kept exchanging responsibilities as each individual UAV continued to attempt further optimizing the current solution.

4.2 Theoretical evaluation (simulating large collectives)

4.2.1 Experimental setup & data collection

The implementation is a re-coding of the original software running on the UAVs, running on a single machine (PC). Each cell is simulated individually and in series so no communication protocols and hand-over mechanisms needed to be implemented. This makes possible the simulation of large colonies and very large numbers of signals in a timely fashion. In addition, no actual signal processing is performed, the measured values are simply reflection of the processing power allocated to each signal.

Cells are activated in a random order and randomly select a signal and a suitable neighbouring cell to exchange the signal with. The resulting signal allocation and the

changes in focus are recorded after $400 / 10^4$ (colony size) cell activations (i.e., after all cells have been activated once). The changes of one such cycle are aggregated into one data point).

Our cells receive outside (simulated) stimuli in the form of desired focus levels for specific signals. Cells calculate a performance value indicating how far they differ from the focus of all signals currently processed by them and attempt to optimise this (i.e., reduce the offset to zero). In the absence of such triggers, cells will aim to optimise the provided signal quality (again, over all signals processed by them) uniformly.

When attempting re-allocation of signals between two cells the values used are the actual value for the cell currently processing the signal and a hypothetical value (i.e., the value *as if the cell was processing this signal*) for the other cell.

The overall signal processing quality across the colony as well as (where applicable) the sum of the discrepancies between provided and desired signal quality are recorded.

We initially simulated 10^4 cells, arranged homogeneously so that each cell is connected to 6 neighbours. Collectively these cells were covering 10^6 signals and in the absence of any triggers, the cells could theoretically settle into a perfect signal-to-cell allocation. We were interested in two behaviours of the population: (1) how the system would behave in the absence of triggers (convergence) and (2) how well a population in a stable configuration could collectively re-allocate signal ownership to meet temporary triggers. Since the latter impacts a small number of cells only, we repeated the experiments with a smaller population to show behaviour of individual cells as graphs. For that smaller experiment we simulated 400 cells at the same signal/cell ratio (i.e., with 4×10^4 signals). The results for the simulation with 400 cells and for the simulation with 10^4 cells were virtually identical.

4.2.2 Discussion (convergence properties)

As the frequency distribution of focus levels across the population of cells (shown in Figure 3) shows, the cells start out with very wide focus but the system quickly converges towards the (known) optimal configuration. There is no noise in the system, all cells are identical and the topology and cell/signal ratio is designed to allow for an optimal state. While this is not realistic in the application sense, it shows that the system does perform as predicted.

We predicted that for larger collectives of cells the convergence would be, if anything, even smoother as the fluctuations caused by the stochastic choice are dampened for

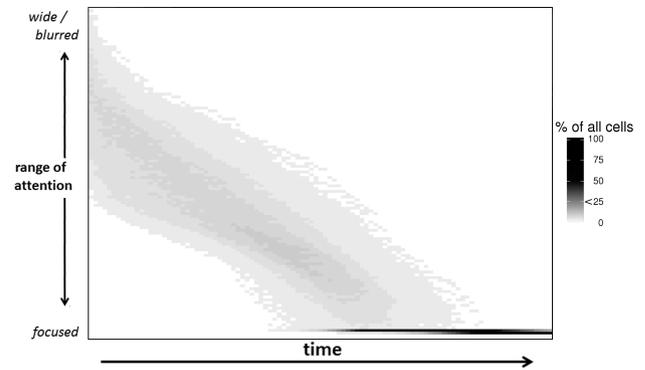


Figure 3: Shown is the frequency distribution of observed focus levels (y-axis, going from *extreme wide focus*, top, to *highly focused*, bottom) over time (x-axis). The results shown were obtained from simulating 100×100 cells arranged in a perfect hexagonal mesh, each cell with 6 neighbours. Collectively, these cells processed 10^6 signals, ≈ 100 per cell. Due to the set-up of the simulation the best possible solution is known: complete coverage at optimal focus levels can be achieved with all cells operating at the same (highly focused) level. The allocation of signals to cells at start time is random and cells are initialised to a wide focus level. The system quickly and consistently settles into the optimal configuration.

increasing colony sizes (*law of big numbers*). And indeed, we found that increasing the colony size does not affect the outcome (when measured in the time it takes to settle into the known optimal solution). This is unsurprising as each cell only interacts with its 6 neighbours. Due to the stochastic nature of the approach, some fluctuations are intended to happen even after a steady state has been reached. Therefore, in this simulation the colony size had no impact on the performance of the individual cell. And indeed, the results obtained from larger collectives of cells are almost identical to the one obtained from the smaller simulation of 400 cells and 4×10^4 signals.

4.2.3 Discussion (self-organization)

In a second investigation some selected cells are subjected to triggers (specific signals are set to require a bigger share of the cell's processing power). The cell population can address this problem by increasing the focus of some cells. This has to be off-set by neighbouring cells widening their focus to cover signals dropped by the now focused cells. So as to be able to investigate whether this happens, we only introduced these triggers to the system after a population had reached steady state (i.e., after the right most point of the convergence shown in Figure 3 was reached). Figure 4 shows the resulting change in focus levels for this.

We simulated only 400 cells and introduced 4 triggers. These values were chosen to generate meaningful im-

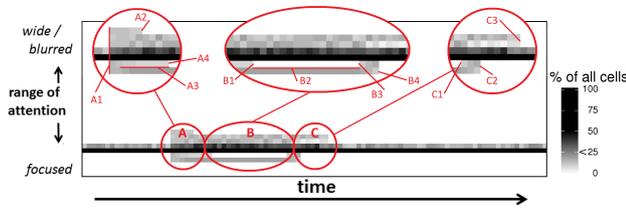


Figure 4: The reaction of the cell population to triggers after the population has reached the steady state shown at the far right end of Figure 3. Shown is the frequency distribution of observed focus levels (y-axis, going from *extreme wide focus*, top, to *highly focused*, bottom) over time (x-axis). The results shown were obtained from simulating 400 cells. The cell population starts out converged to a stable condition. Triggers are introduced in A and focus requirements are met. Localised increased focus is maintained throughout B until the triggers are removed in C, allowing the population to return to the initial stable condition. In detail: after the introduction of the triggers a portion of the cells either increase their focus or widen it (A_1). Some cells widen their focus more than strictly necessary but this is optimised and by A_2 their focus is increased again. Similarly, too many cells react to the triggers by increasing their focus. In this case it takes a bit longer for the cells to optimise (A_3). At A_4/B_1 a stable (and near optimal) configuration is reached. For the duration of B_2 the focused cells do not change. After the triggers are removed at B_3/C_1 the focused cells gradually widen their focus again until B_4/C_2 . As before and because focusing takes longer than unfocusing, it takes the cells with wider focus a bit longer to settle back into the stable initial condition (C_3).

ages (cf. Figure 4) but the performance for larger systems is for all intents and purposes identical. The choice for a smaller simulation is justified by the consideration that a frequency distribution such as the one provided is only meaningful if a significant number of cells do change. We compared the results to both, the affected sub-set of cells in the simulation with 10^4 cells as well as to the full population when 100 triggers were introduced (evenly spaced so as to not interfere with each others impacts).

Figure 4 highlights three stages of this experiment:

- A The moment when the triggers are introduced. This requires some cells to increase their focus, which in turn causes some of their neighbours to take on those signals that cannot be covered anymore.
- B The period when the population has settled into a new steady state and
- C the stretch of time after the triggers are removed, allowing the population to return to the stable condition the system was in at the beginning of the experiment.

Evidently it requires more time to focus attention than it does to reduce focus. This is explained by the fact that reducing coverage is much harder to achieve (it is only possible when all signals in the outer rim of coverage have been handed over to a neighbour) than the widening thereof.

Furthermore, the provided heatmap nicely shows the continuing fluctuation in the population of cells. This is expected (due to the stochastic nature of the approach) and intended (it constitutes a continuous exploration of the solution space).

4.2.4 Discussion (our implementation and model)

Our simulations are set in 2D space; we simplify the coverage to be circular. This means that changes in coverage (i.e. changes in the radius r) affect the attention for each covered signal (i.e., the fixed processing power evenly distributed over the covered area A) in a non-linear way: $A = \pi r^2$. In a 3D space this is worse ($A = 4\pi r^2$) and this increases with more dimensions. Our assumption of circular coverage areas does not cause this, it merely makes the computations for our simulations faster. We argue that this non-linearity is posing a computationally hard problem, a problem we have previously solved in various settings with the proposed approach ([10, 12, 19]).

We furthermore acknowledge that our simulation focuses on allocating signals to exactly one cell, something which, in the context of biological examples, is somewhat contrived. In the real world, cells would likely use all signals in their coverage even if some other cell also processed them. The simulations were meant to show the ability to contract coverage and the implemented scenario was tailored to show that. We argue that these simplifications do not impact the claim that our approach lends itself for use in extremely simple computational nodes.

5 Conclusion

We present a self-organising sensing system of a population of nodes, which are collectively processing a large number of signals. Our approach is the application of a method which we successfully applied to small numbers of complex devices (UAVs) which were required to collaborate so as to cover, collectively as a swarm, an area and e.g., provide video feeds at varying resolution.

It has long since been our intention to use the underlying method for large collectives of simple units to enable them to self-organise their attention of focus and processing powers to raw signals from rudimentary sensing nodes. In a way, this requires the opposite of what the existing and validated implementation to UAVs demands with regard to computational power and complexity. The idea is to enable each of these simple units to self-

determine which of the signals processed by them should be receiving increased focus, or, in other words, when to reduce the width of their sensing range.

The inspiration for this comes from work in cognitive psychology (cf. [4]) on how very large numbers of individual signals received e.g., by photo-receptors in the eye are aggregated - through a number of layers of filters - into highly processed (and biased) data that is eventually processed at a conscious level. In this case, the sheer amount of visual input perceived continuously requires selective processing [5] and various levels of classification before some are selected for additional (cognitive) resource allocation [22]. The suggested approach, when implemented across large populations of units, could serve as a tunable facilitator for such selection processes.

We do not intend to provide alternative models to e.g., [3], [2] on how attention is focused in living animals, nor have we used a realistic implementation of how signals are perceived in e.g., the eye. Our intention is to advocate further work into designing and building large colonies of simple sensing nodes equipped with minimal signal processing capability and to use the proposed approach in low level layers of filters for them.

The literature tells us that the evolution of sensory systems in biology are driven by (a) the task, (b) the limitations of the individual nodes and (c) the environment (embodiment) of the colony [23]. On a higher level, this equally holds for robots; and indeed, the application of nature inspired approaches such as self-organization and emergence to low level signal processing is increasingly considered in the field [14].

Therefore, and in the context of this special issue, we would like to propose our approach - which has been tried and tested for small numbers of complex devices - for consideration to the community. It is straight-forward to see how the allocation mechanism can be driven by simple mathematical operations, though the details of this will entirely depend on the specific hardware (bio-ware?) and application. In the years to come we expect to see more and more opportunities for environment-tailored and task-specific usage of the presented idea.

References

- [1] A. Bandura, Social cognitive theory: an agentic perspective, *Annual Review of Psychology*, 2001, 52(1), 1–26
- [2] R. Desimone, J. Duncan, Neural mechanisms of selective visual attention, *Annual Review of Neuroscience*, 1995, 18, 193–222
- [3] P. R. Roelfsema, Cortical algorithms for perceptual grouping, *Annual Review of Neuroscience*, 2006, 29(1), 203–227
- [4] J. R. Anderson, *Cognitive Psychology and Its Implications*, A series of books in psychology, W. H. Freeman, 1995
- [5] K. McAlonan, J. Cavanaugh, R. H. Wurtz, Guarding the gateway to cortex with attention in visual thalamus, *Nature*, 2008, 456(7220), 391–394
- [6] P. R. Roelfsema, Attention – voluntary control of brain cells, *Science*, 2011, 332(6037), 1512–1513
- [7] J. P. Gottlieb, M. Kusunoki, M. E. Goldberg, The representation of visual salience in monkey parietal cortex, *Nature*, 1998, 391(6666), 481–484
- [8] J.-L. Deneubourg, Application de l'ordre par fluctuations a la description de certaines étapes de la construction du nid chez les termites, *Insectes Sociaux*, 1977, 24(2), 117–130
- [9] Patent Application, Method and system for providing data from a plurality of sensing devices, PCT/EP2015/062810 (pending), H. Hildmann, S. Nicolas, M. Martin (inventors), Dec. 2016
- [10] M. Almeida, H. Hildmann, G. Solmaz, Distributed UAV-swarmbased real-time geomatic data collection under dynamically changing resolution requirements, In: *UAV-g 2017 – International Conference on Unmanned Aerial Vehicles in Geomatics*, Bonn, Germany, Sept. 2017
- [11] S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, E. Bonabeau, *Self-Organization in Biological Systems*, Princeton University Press, 2001
- [12] H. Hildmann, S. Nicolas, F. Saffre, A bio-inspired resource-saving approach to dynamic client-server association, *IEEE Intelligent Systems*, 2012, 27(6), 17–25
- [13] J. Bartholdi, D. Eisenstein, A production line that balances itself, *Operations Research*, 1996, 44(1), 21–34
- [14] R. Pfeifer, M. Lungarella, F. Iida, Self-organization, embodiment, and biologically inspired robotics, *Science*, 2007, 318(5853), 1088–1093
- [15] T. D. Seeley, *Honeybee Democracy*, Princeton University Press, 2010
- [16] R. Beckers, O. Holland, J.-L. Deneubourg, From local actions to global tasks: stigmergy and collective Robots, In: *Proceedings of the Workshop on Artificial Life*, MIT Press, 1994, 181–189
- [17] J.-L. Deneubourg, S. Aron, S. Goss, J. M. Pasteels, The self-organizing exploratory pattern of the Argentine ant, *Journal of Insect Behavior*, 1990, 3(2), 159–168
- [18] E. Bonabeau, M. Dorigo, G. Theraulaz, Inspiration for optimization from social insect behaviour, *Nature*, 2000, 406, 39–42
- [19] H. Hildmann, M. Martin, Adaptive scheduling in dynamic environments, In: *Proceedings of the 2014 Federated Conference on Computer Science and Information Systems*, *Annals of Computer Science and Information Systems*, 2014, 2, 1331–1336
- [20] R. Altawy, A. M. Youssef, Security, privacy, and safety aspects of civilian drones: A survey, *ACM Transactions on Cyber-Physical Systems*, 2016, 1(2), Article No. 7
- [21] F. Saffre, H. Hildmann, J.-L. Deneubourg, Can individual heterogeneity influence self-organised patterns in the termite nest construction model?, *Swarm Intelligence*, 2018, 12(2), 101–110
- [22] M. Cerf, N. Thiruvengadam, F. Mormann, A. Kraskov, R. Q. Quiroga, C. Koch, I. Fried, On-line, voluntary control of human temporal lobe neurons, *Nature*, 2010, 467(7319), 1104–1108
- [23] E. P. Simoncelli, B. A. Olshausen, Natural image statistics and neural representation, *Annual Review of Neuroscience*, 2001, 24, 1193–1216