

Collective Homeostasis and Time-resolved Models of Self-organised Task Allocation

Bernd Meyer
Faculty of IT
Monash University
Melbourne, Australia
bernd.meyer@acm.org

Rui Chen
Faculty of IT
Monash University
Melbourne, Australia
rui.chen@monash.edu

Anja Weidenmüller
Department of Neurobiology
University of Konstanz
Konstanz, Germany
anja.weidenmueller@uni-
konstanz.de

Julián García
Faculty of IT
Monash University
Melbourne, Australia
julian.garcia@monash.edu

ABSTRACT

One of the main factors behind the amazing ecological success of social insects is their ability to flexibly allocate the colony's workforce to all the different tasks it has to address. Insights into the self-organised task allocation methods used for this have given rise to the design of an important class of bio-inspired algorithms for network control, industrial optimisation, and other applications. The most widely used class of models for self-organised task allocation, which also forms the core of these algorithms, are the so-called response threshold models.

We revisit response threshold models with new experiments using temperature regulation in bumblebee colonies as the model system. We show that standard response threshold models do not fit our experiments and present a new, alternative behavioural model. This captures a fine-grained, time resolved picture of task engagement, which enables us to investigate task allocation with a different set of statistical methods (survival analysis). Using these we show that our model fits the experiment well and explains its salient aspects.

We compare the effectiveness of the behaviour that our model describes with that of response threshold models and demonstrate that it can lead to more efficient task management when demands fluctuate. Our results have the potential to provide a basis for the design of new, more efficient task allocation algorithms for dynamic environments and to elucidate important biological questions, such as the functional role of inter-individual variation.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Intelligent Agents; I.6 [Simulation and Modelling]: Applications, Miscellaneous

General Terms

Theory

Keywords

task allocation, division of labour, bio-inspired algorithms, social insects, *Bombus terrestris*

1. INTRODUCTION

The ability of social insects to flexibly adapt their colony organisation to changing environments is a key factor in ensuring their overwhelming ecological success [30, 21]. Underlying this so-called colony plasticity are flexible, self-organised mechanisms of task allocation which ensure that all the different tasks a colony has to tackle simultaneously are serviced according to their importance and urgency, from defence and foraging to nest maintenance and brood rearing. Colony plasticity is a fascinating puzzle, because organisation emerges, without any central coordination, from decentralised interactions between myriads of individuals.

Owing to its flexibility and robustness, self-organised collective behaviour has inspired a wide range of engineering solutions, from control theory [31] to combinatorial optimisation [11]. Task allocation mechanisms, in particular, have been used for applications in factory floor scheduling [6], network control [12], and others.

Despite its central interest in biology and bio-inspired engineering, the individual-based mechanisms of task allocation in insect colonies are only partly understood. While there had been rapid progress in deciphering these mechanisms initially [33, 3], further progress has been hampered by theoretical as much as by experimental challenges:

1. Established mathematical and computational models [4, 35, 13, 3, 14, 26] do not capture a fine-grained, time-resolved picture of individual actions and interactions.

However, we do know from the theory of stochastic processes [18] as well as from studies in insect foraging [16] that timing is a vital ingredient in the analysis of self-reinforced control mechanisms. We may thus expect that it also plays a vital (and so far overlooked) role in task allocation.

2. Established experimental methods [17, 37, 15, 22] do not allow sufficiently precise and flexible control of task-associated stimuli.

Our current work aims to address both issues. We present a new stochastic model for task allocation that integrates timing as a fundamental ingredient. We validate this model with experimental studies of bumblebees *Bombus terrestris* that employ new experimental assays with a fine-grained dynamic stimulus control [39]. We use temperature control as a model task to investigate task allocation, specifically wing fanning for brood cooling. To the best of our knowledge, our work constitutes the first successful attempt to disentangle the effects of temperature stimuli and previous task engagement on the activity of individuals in a temperature control task. Our results challenge the established view that stimulus intensity modulates individual behaviour in a continuous fashion. Rather, individual workers appear to use the stimulus intensity only as a binary indicator to decide whether the task is required at all. We demonstrate that this behaviour makes sense for a colony when the emphasis is on achieving reliable collective homeostasis, as in nature, rather than on counteracting extreme temperatures, as in some experimental investigations [22, 39]. These results have important implications for the design of bio-inspired task allocation mechanisms and shed new light on fundamental biological questions, such as the role of inter-individual variability in insect colonies.

2. BIOLOGICAL BACKGROUND

Different mechanisms have been suggested to underlie task allocation in social insects (reviewed in [3]). Task allocation can, for example, be based on worker polymorphism (task repertoire correlates with differences in morphology) or temporal polyethism (task repertoire changes with age). It is generally assumed that task allocation at the colony level is implemented by behavioural rules at the individual level, i.e. at the level of workers who have no global representation of the colony’s current needs. How the flexibility in task allocation emerges from individual worker behaviour remains a challenging question. What set of behavioural rules does a worker follow for choosing her occupation, which then produces the required group level outcome?

Most empirical studies investigating inter-individual differences have focused on the genetic basis for differences in task choice among workers [24]. Other empirical data is scarce. Previous experimental work that investigated temperature control tasks has only been performed on temperature ramps [29, 22, 39]. On a ramp temperature is a function of time. For a process that is not time-homogeneous, it is thus impossible to reliably tell apart the influences of time and of temperature. The assumption that time does not have an influence must not be made a priori as it can (and will), if made incorrectly, lead to misconceptions. For example, the observation that workers start to engage at different temperatures has commonly been interpreted as evidence that workers decide on task engagement based on

temperature thresholds. However, with a constantly rising (or falling) temperature this could simply be caused by different delays or reactions times. To distinguish such effects, we conduct experiments with different constant temperature settings.

In our experiments we exposed individual bumblebees from young colonies to brood dummies at different temperatures and logged their fanning activity in 1 second intervals over a 30 minute period for each individual. 32 workers from two colonies were tested, 16 at 32°C and 42°C, another 16 at 32°C and 47°C. The setup consisted of a test arena into which a single temperature-controlled brood dummy protruded centrally (Figure 1). At the beginning of each test, the test arena offered a 32°C brood dummy, reflecting optimum brood temperature [37]. For each experiment, one worker was picked up gently by her wings from the brood area of her colony under red light using forceps and placed directly on the brood dummy in the test chamber. The test chamber was closed with a Plexiglas lid and the bee was left undisturbed for a 10 min acclimatisation phase. Her fanning behaviour was then continuously observed and recorded during 10 min. Next, the test arena containing the bee was very carefully lifted from the brood dummy and placed on the second, adjacent brood dummy, which now protruded centrally into the test arena and had a temperature of either 42°C or 47°C. The fanning behaviour of the bee on this second brood dummy was recorded for another 20 minutes. At the end of each experiment, the bee was marked and returned to her colony.




Figure 1: Test arena setup: (b) brood dummy; (h) heating plate; (i) insulation.

3. RESPONSE THRESHOLD MODELS

The most commonly used model of task allocation is the *probabilistic response threshold model* (PTM, [4]). It asserts that the probability of an individual to engage with a task depends on an individual threshold and the intensity of the stimulus. This response probability p is calculated as:

$$p(t) = \frac{s(t)^2}{s(t)^2 + \Theta(t)^2}$$

where $s(t)$ is the intensity of the stimulus at time t and $\Theta(t)$ the individual’s threshold. The threshold itself is either assumed to be fixed [4] or to adapt with task execution. In particular, the *threshold reinforcement model* [35] assumes that the threshold decreases as a result of engaging with the task and increases by itself over time if the individual does not engage with the task:

$$\Theta(t+1) \leftarrow \begin{cases} \Theta(t) - \xi & \text{if the task is executed at time } t \\ \Theta(t) + \phi & \text{otherwise} \end{cases}$$

where ξ and ϕ are the rates of learning and forgetting, respectively. The stimulus s itself may, of course, change as a result of the task execution.

There are two important implications of the response threshold model that are in principle verifiable through direct observation:

1. The *fixed response threshold model* asserts that for an individual the probability to engage with the task increases with stimulus intensity;
2. The *threshold reinforcement model* additionally dictates that the probability to engage with the task increases with continued/repeated task execution.

3.1 Testing the PTM for Bumblebees

Before proceeding we need to discuss how time enters into the response threshold model. The model is usually implemented with fixed discrete time steps of size Δt [4, 35]. The model implications are straight forward if the stimulus level is constant. In this case, the fixed threshold model implies a fixed response probability and thus a constant average rate λ of response. Assuming that p is scaled to be the probability that an individual engages in $\Delta t = 1$ this rate is simply $\lambda(t) = p(t)$. In the limit for a fine-grained time resolution,

$$\Delta t \rightarrow 0 \text{ with } p/\Delta t = \text{constant},$$

a Poisson process with rate constant $\lambda(t) = p(t)$ arises [18]. The two hypotheses above imply that

1. In both versions of the threshold model, the rate constant λ of the Poisson process must increase with temperature;
2. If we assume reinforced thresholds, the Poisson process is not time-homogeneous. $\lambda(t)$ increases over time as the individual repeatedly engages with the task.

To test these hypotheses we conduct our experiments with different *constant* temperatures $T = 42^\circ C$ and $T = 47^\circ C$. For each individual i we record two sequences f and s capturing the start and end time of each uninterrupted fanning bout with a resolution of 1 second. $f_{i,j}^T$ is the start time of the j -th fanning bout of individual i for a run performed at temperature T and $s_{i,j}^T$ the corresponding end time of this fanning bout. We investigate the intervals between subsequent fanning bouts $w_{i,j}^T = f_{i,j+1}^T - s_{i,j}^T$, which allows us to analyse the dependency of the rate function $\lambda(t)$ on time and temperature.

To test hypothesis no. 1 we analyse the influence of temperature on the rate function of the processes defined by the event times $w^T = \cup_{i,j} w_{i,j}^T$. Specifically, we use Cox regression [27] to test whether the rate functions of w^{42} and w^{47} differ systematically depending on temperature. Note that Cox regression is a non-parametric technique, so that this analysis does not depend on the assumption that the underlying process is indeed Poissonian. The required proportional hazards assumption applies since the Schoenfeld residuals are uncorrelated with the ranked event times (Pearson Correlation 0.03, graphically supported in Figure 2). The relative risk factor for w^{42} and w^{47} is 1.016 with confidence interval [0.993...1.039], which clearly shows that temperature does not significantly influence the rate $\lambda(t)$ of engagement with the task. This contradicts PTM hypothesis no. 1.




Figure 2: The Schoenfeld residuals over the ranks of event times.

Note that this analysis only concerns individuals who engage with the task *at all*. There are a small number of individuals (5 at $42^\circ C$, 2 at $47^\circ C$) who do not show any fanning activity whatsoever. As these do not produce any bout intervals, the analysis is implicitly conditioned on individuals who perform at least one fanning bout.

To test hypothesis no. 2 we cannot use Cox regression since the required proportional hazard assumption cannot be validated. Instead, we use a graphical analysis. We partition the event data $w^T = \cup_i b_i^T$ into strata b_i according to the repeat number i of the bout. b_i^T is the subset of w^T that only contains event times originating from the i -th bout of any individual. Figure 3 shows the log-log survival plots for the event times b_i with a separate curve for three representative samples of b_i . The remaining b_i curves look similar and have been omitted for clarity. If the rate would systematically increase with previous task engagement, we would see the curves for later bout strata b_j consistently above those for earlier bout strata b_i (for $i < j$). This is not the case, and no systematic influence on the repeat number of the bout on the rate is visible. This contradicts PTM hypothesis no. 2.




Figure 3: The log-log survival curves in log-scaled time for b_i , $i = 1, 10, 20$.

We conclude that the PTM is not applicable to our experimental setting. To the best of our knowledge, our experiments are the first to decouple (1) the influence of time and the influence of stimulus intensity on the probability to engage with the task and (2) the probability to engage with the task from the average activity level. While it is possible that the PTM applies to other species and other tasks, it is equally possible that the successful analysis of previous experiments has been foiled by the lumping of these factors.

4. A TIME-RESOLVED MODEL OF TASK ENGAGEMENT

Whether or not to engage in a task, such as fanning, is a decision that is constantly reassessed by the individuals. Our experimental observations show that even at a constant temperature individuals assess the temperature several times before a decision to fan is made. One observable indication of such an assessment is the antennation of the brood dummy. Typically, they fan only for a short while before they stop and re-assess the situation. This is clearly visible in the activity pattern displayed in Figure 6.

The standard response threshold model does not account for these complex timing patterns of individual behaviour. To address this, we present an explicitly time-resolved model of task allocation, or, more precisely, of emergent task distribution. The model defines the individual level of behaviour. Its central difference to existing threshold models is that it explicitly accounts for the ongoing re-assessment of task engagement with a stochastic model. We will investigate how collective homeostasis arises from this model in Section 5.

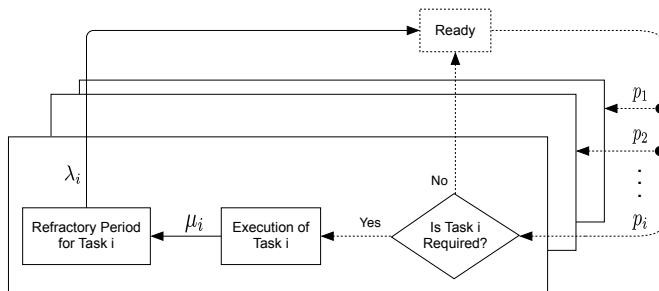


Figure 4: The time-resolved task polling model for a single individual.

We assume that there are N tasks that an individual w can choose to perform at any given time. Figure 4 illustrates the flow of the decision making process. Executing a task is a three step process that involves polling task stimuli. When w is ready to take on a new task, she will first decide which task i to execute according to the task polling probabilities p_i . She will then test whether the task-related stimulus s_i is above the corresponding threshold Θ_i . If so, w will engage in i , otherwise she will poll for another task j according to the probabilities p_j . Once engaged in a task, w will execute this task for a certain amount of time after which she enters a refractory resting period. Abandoning the task as well as finishing resting are stochastic processes that in the absence of other assumptions are best modelled as Poisson processes with rates λ_i and μ_i , respectively. After the completion of one such task cycle the worker is immediately free to engage in another task or to resume the same task as before. The parameters Θ and p may vary between individuals.

Note that we use a deterministic threshold decision. This is for two reasons: (1) As shown in Section 3.1 our experimental data does not support the assumption of a probabilistic threshold; (2) a deterministic threshold is the simplest possible assumption and fully sufficient for homeostasis since the system works as a closed control loop: Each individual assesses the temperature stimulus, takes corrective action for a while if required, and then reassesses the sit-

uation. While this is not a continuous-time process as in classical control theory, it approximates this for sufficiently fine time intervals. As long as a sufficient number of individuals with threshold at the optimal temperature are available, the colony will achieve homeostasis at this temperature.

4.1 Reduced Model for Empirical Studies

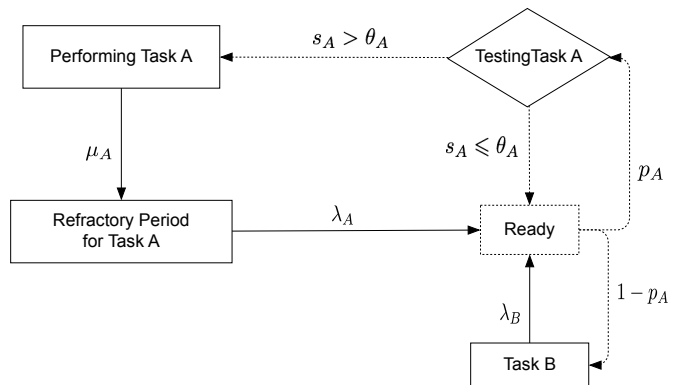


Figure 5: Simplified time-resolved task polling model for a single observable task and for a single individual.

The above model makes the most general assumption that a worker chooses between N different behaviours. However, in an experimental setting it is exceedingly difficult to exactly control multiple task-associated stimuli simultaneously. Thus, most experimental settings control and observe only a single task, such as fanning in our experiments. For the analysis we must, however, account for the fact that other behaviours may be executed in between the controlled tasks. This includes observable behaviours, such as self-grooming and running, as well as possibly unobservable ones.

To apply the model to these experiments we thus reduce it to only distinguish between a known, observable task A and a set of alternative tasks B (see Figure 5). Task $B = \{B_1, \dots, B_{N-1}\}$ is thus a lumped meta-task that represents a mix of possible tasks.

In this context it is interesting to consider that in any social insect colony large proportions of the workforce are typically passive and do not seem to execute any observable task. The proportion of the passive workforce can reach up to 50% of the colony [9, 23], and while it is often assumed that this is the colony's way to provide a reserve workforce for high demand situations, there is no consensus on the functional role, if any, of this passive workforce [7]. While we cannot be certain that a worker who is not executing the observable task A is indeed simply passive, the possibilities for B_i clearly include “not doing anything.”

One specific potential reason for not doing anything is that workers may need to recover from a previously executed task. Our model accounts for this explicitly with refractory periods that are entered immediately after task execution. Note that these could have zero residency time if no refractory period is required.




Figure 6: Fanning activities of bumblebee workers at a constant brood dummy temperature of 42°C . Each row represents a single worker. Each column represents a time interval of 1 second. Fanning activity is represented by the gradient of shading, white boxes denote inactivity or activity other than fanning. As the chart cannot be displayed at full resolution, very short interruptions of activity are only visible in the gradient of the shading rather than as explicit white gaps.

4.2 Empirical Evidence

We implemented a statistically faithful simulation of this model based on the Gillespie next reaction method [19] (see Algorithm 1). The time series data gathered in our experiments is consistent with the data generated by our model.¹

Algorithm 1 Simulation algorithm for the time-resolved task polling model

```

1:  $t := 0$ 
2: initialise all individuals  $w_i$ 
3: while  $t < t_{end}$  do
4:    $H \leftarrow \sum_i r_i$  (* total hazard *)
5:   draw next event time  $\Delta t \sim \text{Exp}(H)$ 
6:    $t \leftarrow t + \Delta t$ 
7:   choose  $w_i$  according to probabilities  $p(w_i) = r_i/H$ 
8:   (* the individual concerned by this event *)
9:   if  $\text{state}(w_i) = A$  then
10:     $\text{state}(w_i) = R_A$ 
11:     $r_i \leftarrow \lambda_A$ 
12:   else
13:     $n \leftarrow \text{random}(0, 1)$ 
14:    if  $n < p_{A,i} \wedge s > \theta_i$  then
15:       $\text{state}(w_i) = A$ 
16:       $r_i \leftarrow \mu_A$ 
17:      (*  $w_i$  executes Task A *)
18:    else
19:       $\text{state}(w_i) = B$ 
20:       $r_i \leftarrow \lambda_B$ 
21:    end if
22:   end if
23: end while

```

4.2.1 Group Behaviour

Figure 7(a) shows the empirical distribution of all fanning intervals $w_{i,j}$. This is clearly not an exponential distribution, as would be expected for a straight-forward Poisson extension of the response threshold model (cf. Section 3.1). Instead, the distribution is power-law like as is recognisable from the straight slope in the log-log scaled histogram.²

¹Model simulations and all statistical evaluations were performed in Mathematica V. 10.2.0.0. Source code is available from the first author.

²As usual with empirical distributions it is difficult to reliably identify an exact power-law. Firstly, any empirical data can only support a truncated power-law. Secondly, it is not easy to clearly distinguish this from a log-normal. Our main point here is not to confirm the exact shape of the distribution, but that it is clearly

Our model reproduces the distribution of the fanning intervals well. Figure 7 compares the empirical distribution to the simulation data of our fitted model and the basic response threshold model (assuming its Poissonian extension). Goodness of fit tests confirm what is visually obvious: the Kullback-Leibler distance is 0.076 between the distributions of empirical data and simulation data and 7.879 between the empirical distribution and the distribution of the Poisson process. This means that an AIC-based model choice must discard the Poisson model [1].

4.2.2 Individual Behaviour

On the individual level, we see a large variation of bout interval characteristics. Figure 8 shows typical cases. Some individuals exhibit a power-law like statistics in their fanning intervals, while in other cases the distribution more closely resembles an exponential distribution. Our model explains this variability.

Figure 9 shows the simulation data for a single individual with varying values of p_A . For low values of p_A , the individual fanning statistics is power-law like whereas for high values of p_A it approximates an exponential distribution. This is explained as follows: B is a lumped meta-task in which each B_i has a different characteristic duration $\frac{1}{\lambda(B_i)}$. The inter-bout intervals are thus a mixture of waiting times caused by the refractory period for A and the alternative tasks B_i .

For high probabilities p_A the mixture consist almost exclusively of the refractory period for A and the inter-bout intervals thus approximate a Poisson process with rate $\lambda(A)$. For low values of p_A the inter-bout intervals are mostly determined by the execution times for B , which itself is a mixture of Poisson processes with different rates, and such a mixture process exhibits a power-law like timing even if each individual process B_i is Poissonian [20].

5. COLLECTIVE HOMEOSTASIS CONTROL

Collectively bumblebee colonies manage to keep the nest temperature remarkably stable at the ideal target temperature of 32°C over very long periods [37]. This can be explained with the closed control loop behaviour at the indi-

not exponential. For simplicity we will refer to this as power-law like in the remainder of the paper.




Figure 7: Distribution of lumped fanning bout intervals on log-log scales for empirical data, simulation data ($N = 25, t = 1400, \lambda_A = 0.8, \mu_A = 0.1, \lambda_B \sim U(0, 1), p_A \sim U(0, 1)$), and the probability density function of the Poisson process fitted to the rate constant of the empirical data.

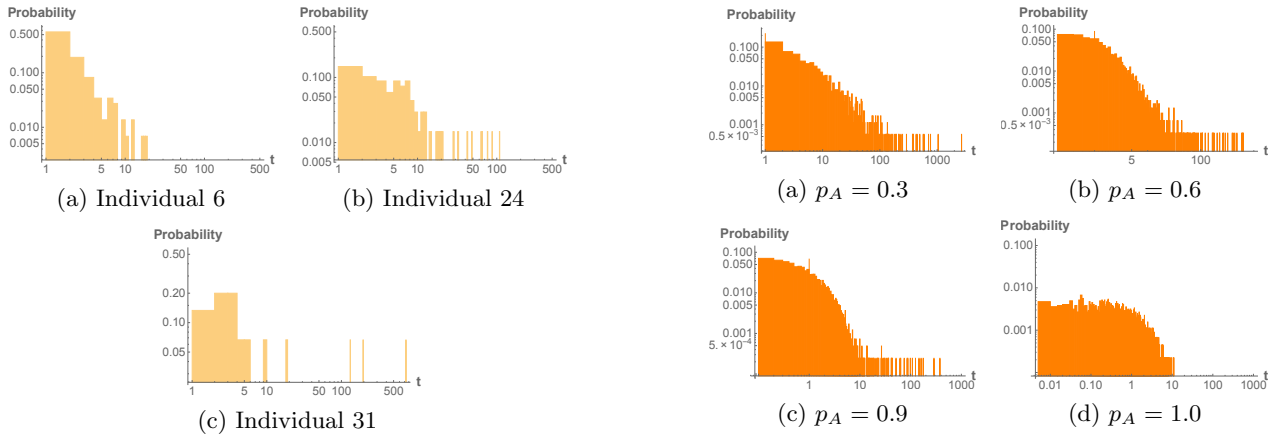


Figure 8: Histogram of inter-fanning times of individual bumblebee workers (experimental data).

Figure 9: Distribution of individual fanning bout intervals on log-log scales for simulation data with $t = 50000, \lambda_A = 0.8, \mu_A = 0.1, \lambda_B \sim U(0, 1), p_A$ varying from 0.3 to 1.

vidual level, even though some social influences may also enter into the process [10].

However, at least two important questions are left open: (1) How quickly can the colony manage to equilibrate the nest temperature or, equivalently, how stable is the nest temperature if the environment conditions fluctuate; and (2) which proportion of the workforce is engaged in the thermoregulation process.

When the temperature is at equilibrium, which in reality it is most of the time, workers drop in and out of the active cooling work force. As some individuals pause to rest, to check the stimulus, or to attend to other tasks, new individuals engage in the process. Thus, the absolute number of workers engaged in the cooling process over extended periods of time may vary significantly, in principle between just the number required to balance the energy input and the whole colony. This raises important questions regarding the efficiency of the process. It is generally assumed that some degree of specialisation in tasks leads to more efficient task execution [21]. A variety of effects enter into this. Intrinsic differences (e.g. morphological) may enable some individuals to perform the task more efficiently [24]. Task execution itself may enable the individual to learn how to execute the task more efficiently [32, 25, 8]. Specialisation may also reduce the amount of overhead required for task allocation. In a coarse simplification, we may assume that it is of advan-

tage to the colony if the number of workers engaged in the cooling task is smaller, thus keeping the remaining workers free to perform other tasks.

5.1 Inter-individual Variability

On the background of this, we can ask whether the wide variability that we observe in the behaviour of individuals ([36], see also Figures 6 and 8) serves a functional role or whether it is just an epiphenomenon.

When temperature and cooling activity are at steady state a constant swapping of individuals into and out of the active cooling work force takes place. If all workers behave identically all have an equal chance of engaging in the swapping process so that a large fraction of the colony participates over time, but each individual only contributes a little. If, on the other hand, there is wider distribution of p_A values those workers with a high value of p_A will on average engage more quickly (and thus more frequently) with the task, effectively preventing workers with a lower p_A to engage.

This is borne out by simulations. We extend our simulation to capture a simple model of the physics of the collective cooling process. The virtual nest is exposed to a given ambient temperature and additionally produces internal heat. Standard Newtonian heat conduction processes are used to approximate the effect on the nest. To simplify, we assume

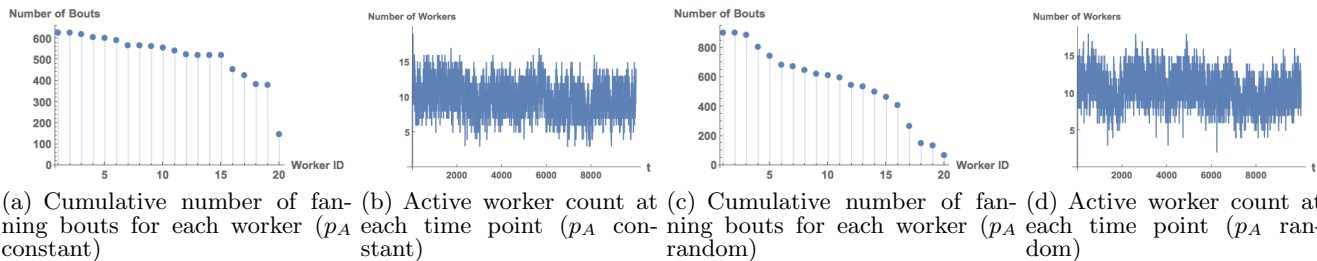


Figure 10: Fanning activity in a colony of identical workers versus a heterogeneous colony. $N = 20$, $\lambda_A = 0.8$, $\mu_A = 0.1$, $\lambda_B \sim U(0, 1)$ and $p_A = 0.5$ for 10(a), 10(b) and $p_A \sim U(0, 1)$ for 10(c), 10(d). Absolute activity for each worker is displayed as the total time during which a worker was actively engaged in fanning. Workers are sorted from left to right in descending order of their p_A values.

that all bees have identical fixed cooling capacity.

The discussed effect can clearly be demonstrated qualitatively (Figure 10). The simulations confirm that the colony manages to keep the temperature stable with slight fluctuations around the target temperature. The absolute average count of workers engaged in fanning at any given point of time is constant within fluctuations, as it must, and it is a function of the individuals’ cooling capacity and the overall energy input.

For a colony of identical workers far more individuals are involved in the cooling process than for a colony where the p_A values are distributed uniformly between 0 and 1. In the homogeneous colony all but one worker (95%) perform between 400 and 600 units of work each, while just 55% perform more than 600 units in the heterogeneous colony. In effect, almost the whole homogeneous colony is significantly involved in the task and only about half of the heterogeneous colony. Other distributions of p_A could clearly lead to more extreme differences giving rise to a more sharply defined division of labour.

Thus, the observed intrinsic variability at the system level may have a functional role in enabling the colony to achieve collective homeostasis in a more efficient way.

5.2 Comparison to Threshold Reinforcement

It is instructive to compare the effect of fixed intrinsic inter-individual variability to the effect that would be achieved by a threshold reinforcement mechanism. While we have shown above that the assumption of threshold reinforcement does not fit with our experimental data, it is entirely possible that some form of adaptive thresholding exists for different species, tasks, or timescales ([36], cf. Section 3.1). More importantly, from the perspective of bio-inspired algorithms concrete biological systems serve as inspiration, but ultimately we are interested in efficient algorithms and control mechanisms. Thus, if threshold reinforcement yields a better control method we should adopt this for engineering purposes, regardless of whether it exists in biological systems or not.

Threshold reinforcement can achieve a specialisation effect for particular tasks without relying on an intrinsic distribution of thresholds or probabilities [35]. When an individual engages in a task it lowers its propensity to engage in the task again by lowering its threshold for this task. This leads to a positive feedback loop: individuals that engage with a task will on average engage with it more frequently in the future thus further increasing their propensity to do so.

This positive feedback is moderated by negative feedback: the threshold has a self-driven dynamics that causes it to rise unless counterbalanced by the lowering due to task execution. Higher propensity to engage in a particular task effectively constitutes a specialisation.

In any such feedback system, the question whether it can adapt to changing conditions depends on the balance between positive and negative reinforcement [28]. If the mechanism specialises very quickly and is robust against changes, it unspecialises slowly (and conversely). This is because, for strong feedback, the positive reinforcement derived from executing a task overwhelms the “forgetting” that is required to unspecialise. Weak positive reinforcement, on the other hand, renders the system unable to adapt quickly. In other words, flexibility and stability are conflicting goals.

How a system should ideally operate depends on the dynamic characteristics of the environment. In a fluctuating environment strong specialisation by reinforcement may not always be useful since significant amounts of the workforce may remain tied up with less important tasks when the demand for an important task rises.

To compare fixed intrinsic variability to specialisation by adaptation, we adopt the threshold reinforcement paradigm in a slightly refined form for our model, where an individual has to choose between the individual task A (fanning) or the lumped meta-task B .

We assume that the probability to choose either task is based on the relative previous experience with executing the two tasks in such a way that the task for which more experience has been gathered is more likely to be executed. The probability $p_{A,i}$ of individual i to poll for task A is thus

$$p_{A,i} = \frac{f(E_{A,i})}{f(E_{A,i}) + f(E_{B,i})}$$

Here $E_{j,i}$ is the experience with task j gathered by individual i and f is a monotonically increasing function.

We assume the individual derives a certain utility from executing a task and accumulates this utility as experience. Experience is gradually forgotten over time. This directly corresponds to a widely used general class of models for *rational choice* [38, 2], a variant of which is frequently used to model decision making in insect foraging [5]. It holds that the probability to choose a particular option is proportional to the *expected* utility for that choice and that this expected utility is the discounted sum of previous experiences. Thus we calculate the cumulative discounted experience of indi-

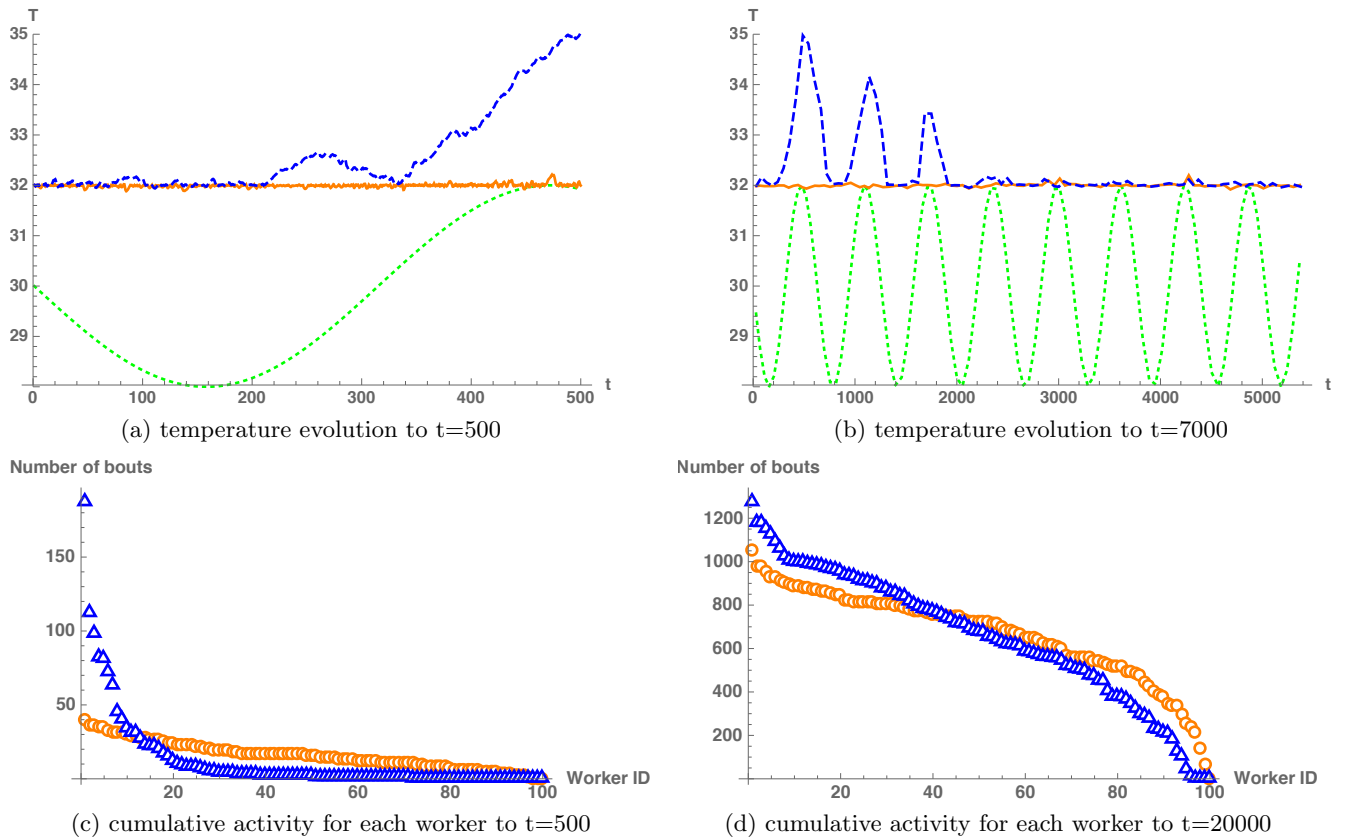


Figure 11: Temperature evolution and total number of fanning bouts for each worker in a colony with $U_i \in \{0, 1\}$ (circles, orange solid line) versus a colony with $U_i \in \{0, 10\}$ (triangles, blue dashed line) for fluctuating ambient temperature (dotted green line). In (c,d) workers are sorted from left to right in descending order of their activity counts.

vidual i with task j as

$$E_{j,i}(n+1) := (1 - \rho) \cdot E_{j,i}(n) + \rho \cdot U_{j,i}(n)$$

where $U_{j,i}(n)$ is the utility that individual i derives from executing task j during its n -th execution of the task. Note that this definition leaves the exact base of adaptation flexible. It could either just be the fact that the individual engages with the task ($U_{j,i} \in \{0, c\}$) or some measure of success of executing the task ($U_{j,i} \in \mathcal{R}$). As we cannot measure the success of executing the task in our experimental setup we shall in the remainder assume that task engagement is the determining factor.

In line with common use in biology [34], we assume $f(x) = (k + x)^\gamma$, where k models a residual probability to choose a task in the absence of experience and γ determines how strongly differences in experience are emphasised.

While a detailed and exhaustive formal analysis is beyond the scope of this paper, we briefly substantiate the above discussion with simulation data. We expose two virtual colonies to fluctuating ambient temperatures that represent the daily or seasonal temperature cycle and record the total amount of time for which each individual performs the cooling task. Colony A is parameterised to achieve fast adaptation ($U_{j,i} \in \{0, 10\}$) while colony B is parameterised for slower adaptation ($U_{j,i} \in \{0, 1\}$). Both colonies use $\rho = 0.05, k = 1, \gamma = 2, \lambda_A = 1, \mu_A = 1, \lambda_B \sim U(0, 1)$.

The simulation bears out the above considerations. Initially the fast adaptation in colony A reduces the number of workers required for fanning and leads to rapid specialisation for $t < 500$ as evidenced by the fact that less than 20% of the workforce perform almost all the work (Fig. 11(c)). However, the colony cannot compensate the rising ambient temperature quickly enough and at first remains adapted to the initially low demand for cooling ($350 \leq t \leq 500$, Fig. 11(a)). Eventually the colony adapts to the environmental fluctuations ($t > 2000$, Fig. 11(b)) but in the process loses the specialisation (Fig. 11(d)). The slowly adapting colony B , on the other hand, copes well with the changes in demand from the start but specialisation never arises. Effectively the adaptation is rendered useless by the fluctuations. In such an environment a system that employs fixed intrinsic inter-individual variability may actually be at an advantage over a system that uses adaptive task preferences. If the environment fluctuates faster than the characteristic time of the adaptation, no persistent task preferences can arise and it is essentially left to chance which individuals engage in a given task. With fixed intrinsic thresholds/probabilities, on the other hand, task preferences are to some degree persistent, i.e. the same workers will tend to perform the same tasks. This may have benefits through long-term learning or if the morphological differences result in different task efficiencies.

6. CONCLUSION AND FUTURE WORK

We have presented a model that captures a fine-grained, time-resolved picture of self-organised task allocation in bumblebees. To test our model we have performed experiments with bumblebees as the model organism and nest temperature regulation as the model task. Our experiments are designed to distinguish the effects of time, repeated task execution and stimulus intensity, which cannot be isolated in previous experiments.

Since the objective of temperature regulation in an insect colony is to achieve homeostasis, our new model is based on a simple closed control loop perspective, in which individuals use stimulus intensity only to judge whether task execution (cooling) is required at all. We have shown that traditional response threshold models are not sufficient to explain the individual behaviour observed in our experiments. Instead, our simple control loop model provides a better fit with our experiment and explains its salient aspects.

The analysis of our model has revealed some important implications for fundamental biological questions. It shows how in dynamic environments fixed intrinsic inter-individual variability can enable a more effective division of labour than adaptation by threshold reinforcement. Ultimately, this allows a colony to reduce the amount of workers tied up with temperature regulation and thus enables it to achieve its tasks more efficiently. We have substantiated this by simulating the temperature control task in a variable environment.

Our findings have potentially important consequences for the design of bio-inspired control algorithms. For applications with dynamic, fluctuating task demands it may be preferable to base the algorithm design on the behaviour described by our model and to utilise fixed intrinsic variability between task processing agents rather than relying on adaptive threshold reinforcement. Future work should explore this from a realistic application-centric perspective.

Some important questions remain open. We have studied variability only in the most fundamental factor, the task polling probabilities p_A . Other factors, such as fanning durations, may also vary between individuals. Likewise, while the task polling probabilities in our experiments are not modulated by stimulus intensity, it is possible that other factors are. The implications of this need to be investigated experimentally and theoretically and we are planning to address this in future work. We also aim at a mathematical analysis of the role of inter-individual variability rather than relying merely on simulations.

7. REFERENCES

- [1] H. Akaike. A new look at the statistical model identification. *IEEE Trans. Autom. Control*, 19(4):716–723, 1974.
- [2] S. Anderson, A. de Palma, and J. Thisse. *Discrete Choice Theory of Product Differentiation*. The MIT Press, Cambridge, MA, 1992.
- [3] S. N. Beshers and J. H. Fewell. Models of division of labor in social insects. *Annu. Rev. Entomol.*, 46:413–440, 2001.
- [4] E. Bonabeau, G. Theraulaz, and J.-L. Deneubourg. Quantitative study of the fixed threshold model for the regulation of division of labour in insect societies. *Proc. R. Soc. Lond. B*, 263(1376):1565–1569, Nov. 1996.
- [5] S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-organization in biological systems*. Princeton University Press, Princeton, NJ, 2001.
- [6] M. Campos, E. Bonabeau, G. Theraulaz, and J.-L. Deneubourg. Dynamic scheduling and division of labor in social insects. *Adpt. Behav.*, 8(2):83–95, Mar. 2000.
- [7] D. Charbonneau and A. Dornhaus. Workers specialized on inactivity: Behavioral consistency of inactive workers and their role in task allocation. *Behav. Ecol. Sociobiol.*, 2015. Published online.
- [8] L. Chittka and H. Muller. Learning, specialization, efficiency and task allocation in social insects. *Commun. Integr. Biol.*, 2(2):151–154, 2009.
- [9] B. J. Cole. The social behavior of *Leptothorax allardycei* (Hymenoptera, Formicidae): time budgets and the evolution of worker reproduction. *Behav. Ecol. Sociobiol.*, 18(3):165–173, Jan. 1986.
- [10] C. N. Cook and M. D. Breed. Social context influences the initiation and threshold of thermoregulatory behaviour in honeybees. *Anim. Behav.*, 86(2):323–329, Aug. 2013.
- [11] M. Dorigo and T. Stützle. Ant colony optimization: overview and recent advances. In M. Gendreau and J.-Y. Potvin, editors, *Handbook of Metaheuristics*, pages 227–263. Springer, New York, 2010.
- [12] F. Dressler. *Self-organization in sensor and actor networks*. Wiley, Chichester, England, 2007.
- [13] A. Duarte, I. Pen, L. Keller, and F. J. Weissing. Evolution of self-organized division of labor in a response threshold model. *Behav. Ecol. Sociobiol.*, 66(6):947–957, Mar. 2012.
- [14] A. Duarte, F. J. Weissing, I. Pen, and L. Keller. An evolutionary perspective on self-organized division of labor in social insects. *Annu. Rev. Ecol. Evol. Syst.*, 42:91–110, 2011.
- [15] N. Duong and A. Dornhaus. Ventilation response thresholds do not change with age or self-reinforcement in workers of the bumble bee *Bombus impatiens*. *Insect. Soc.*, 59(1):25–32, Feb. 2012.
- [16] A. Dussutour, M. Beekman, S. C. Nicolis, and B. Meyer. Noise improves collective decision-making by ants in dynamic environments. *Proc. R. Soc. B*, 276(1677):4353–4361, Dec. 2009.
- [17] J. H. Fewell and S. M. Bertram. Division of labor in a dynamic environment: response by honeybees (*Apis mellifera*) to graded changes in colony pollen stores. *Behav. Ecol. Sociobiol.*, 46(3):171–179, 1999.
- [18] C. W. Gardiner. *Handbook of Stochastic Methods for Physics, Chemistry, and the Natural Sciences*. Springer, 3rd edition, 2004.
- [19] D. Gillespie. *Markov Processes*. Academic Press, San Diego, CA, 1992.
- [20] C. A. Hidalgo R. Conditions for the emergence of scaling in the inter-event time of uncorrelated and seasonal systems. *Physica A*, 369(2):877–883, Sept. 2006.
- [21] B. Hölldobler and E. O. Wilson. *The Ants*. Belknap Press of Harvard University Press, Cambridge, Massachusetts, 1990.

- [22] J. M. Jandt and A. Dornhaus. Bumblebee response thresholds and body size: does worker diversity increase colony performance? *Anim. Behav.*, 87:97–106, 2014.
- [23] J. M. Jandt, N. S. Robins, R. E. Moore, and A. Dornhaus. Individual bumblebees vary in response to disturbance: a test of the defensive reserve hypothesis. *Insect. Soc.*, 59(3):313–321, Aug. 2012.
- [24] R. Jeanson and A. Weidenmüller. Interindividual variability in social insects - proximate causes and ultimate consequences. *Biol. Rev. Camb. Philos. Soc.*, 89(3):671–687, Aug. 2014.
- [25] E. A. Langridge, A. B. Sendova-Franks, and N. R. Franks. How experienced individuals contribute to an improvement in collective performance in ants. *Behav. Ecol. Sociobiol.*, 62(3):447–456, 2008.
- [26] P. Lichocki, D. Tarapore, L. Keller, and D. Floreano. Neural networks as mechanisms to regulate division of labor. *Am. Nat.*, 179, Mar. 2012.
- [27] X. Liu. *Survival Analysis*. Wiley, Chichester, England, 2012.
- [28] G. Nicolis. *Introduction to Nonlinear Science*. Cambridge University Press, 1995.
- [29] S. O’Donnell and R. L. Foster. Thresholds of response in nest thermoregulation by worker bumble bees, *Bombus bifarius nearcticus* (Hymenoptera: Apidae). *Ethology*, 107(5):387–399, May 2001.
- [30] G. F. Oster and E. O. Wilson. *Caste and Ecology in the Social Insects*. Princeton University Press, 1978.
- [31] K. M. Passino. *Biomimicry for Optimization, Control, and Automation*. Springer-Verlag, London, 1st edition, 2005.
- [32] F. Ravary, E. Lecoutey, G. Kaminski, N. Châline, and P. Jaisson. Individual experience alone can generate lasting division of labor in ants. *Curr. Biol.*, 17(15):1308–1312, Aug. 2007.
- [33] G. E. Robinson. Regulation of division of labor in insect societies. *Annu. Rev. Entomol.*, 37(1):637–665, Jan. 1992.
- [34] D. Sumpter. *Collective Animal Behaviour*. Princeton University Press, Princeton, NJ, 2010.
- [35] G. Theraulaz, E. Bonabeau, and J.-L. Deneubourg. Response threshold reinforcements and division of labour in insect societies. *Proc. R. Soc. Lond. B*, 265(1393):327–332, Feb. 1998.
- [36] A. Weidenmüller. The control of nest climate in bumblebee (*Bombus terrestris*) colonies: interindividual variability and self reinforcement in fanning response. *Behav. Ecol.*, 15(1):120–128, Jan. 2004.
- [37] A. Weidenmüller, C. Kleineidam, and J. Tautz. Collective control of nest climate parameters in bumblebee colonies. *Anim. Behav.*, 63(6):1065–1071, June 2002.
- [38] G. Weissbuch, A. Kirman, and D. Herreiner. Market organisation and trading relationships. *Econ. J.*, 110:411–436, 2000.
- [39] C. Westhus, C. J. Kleineidam, F. Roces, and A. Weidenmüller. Behavioural plasticity in the fanning response of bumblebee workers: impact of experience and rate of temperature change. *Anim. Behav.*,