



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/103126/>

Version: Accepted Version

Proceedings Paper:

Bishop, J. M., Martin, Andrew O. and Robinson, Elva Joan Hilda (2016) Local termination criteria for Stochastic Diffusion Search: a comparison with the behaviour of ant nest-site selection. In: Nguyen, Ngoc-Thanh, Manolopoulos, Yannis, Iliadis, Lazaros and Trawiński, Bogdan, (eds.) Lecture Notes in Artificial Intelligence: Proceedings of the 8th International Conference on Computational Collective Intelligence (ICCCI 2016), Halkidiki, Greece. Lecture Notes in Artificial Intelligence. Springer.

<https://doi.org/10.1007/978-3-319-45243-2>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Local termination criteria for Stochastic Diffusion Search: a comparison with the behaviour of ant nest-site selection

J. Mark Bishop¹, Andrew O. Martin¹, and Elva J. H. Robinson²

¹ TCIDA, Goldsmiths, University of London
New Cross, London, UK.

{mark.bishop, andrew.martin}@tungsten-network.com
<http://www.tungsten-network.com/tcida/>

² York Centre for Complex Systems Analysis and Department of Biology
University of York, York, UK.

elva.robinson@york.ac.uk

<https://www.york.ac.uk/biology/research/ecology-evolution/elva-robinson/>

Abstract. Population based decision mechanisms employed by many Swarm Intelligence methods can suffer poor convergence resulting in ill-defined halting criteria and loss of the best solution. Conversely, as a result of its resource allocation mechanism, the solutions found by Stochastic Diffusion Search enjoy excellent stability. Previous implementations of SDS have deployed complex stopping criteria derived from global properties of the agent population; this paper examines two new *local* SDS halting criteria and compares their performance with ‘quorum sensing’ - a natural termination criterion deployed in nature by some species of tandem-running ants. We empirically demonstrate that local termination criteria are almost as robust as the classical SDS termination criteria, whilst the average time taken to reach a decision is around three times faster.

Keywords: Collective Decision Making; Ant Nest Selection; Stochastic Diffusion Search; Swarm Intelligence; Global Search.

1 Introduction

In recent years there has been growing interest in swarm intelligence, a distributed mode of computation utilising interaction between simple agents [21]. Such systems have often been inspired by observing interactions between social insects: ants, bees, termites (cf. Ant Algorithms and Particle Swarm Optimisers) see Bonabeau [10] for a comprehensive review. Swarm Intelligence algorithms also include methods inspired by natural evolution such as Genetic Algorithms [18] [20] or indeed Evolutionary Algorithms [5]. The problem solving ability of Swarm Intelligence methods emerges from positive feedback reinforcing potentially good solutions and the spatial/temporal characteristics of their agent interactions.

Independently of these algorithms, Stochastic Diffusion Search (SDS), was first described in 1989 as a population-based, pattern-matching algorithm [7]. Unlike stigmergic communication employed in Ant Algorithms, which is based on modification of the physical properties of a simulated environment, SDS uses a form of direct communication between the agents similar to the tandem running mechanism employed by some species of ants (e.g. *Temnothorax* species, [14]).

SDS is an efficient probabilistic multi-agent global search, optimisation and decision making technique [23] that has been applied to diverse problems such as site selection for wireless networks [41], mobile robot self-localisation [6], object recognition [8] and text search [7]. Additionally, a hybrid SDS and n-tuple RAM [1] technique has been used to track facial features in video sequences [8] [19]. Previous analysis of SDS has investigated its global convergence [26], linear time complexity [27] and resource allocation [25] under a variety of search conditions. For a recent review of the theoretical foundations, and applications of SDS see Al-Rifaie and Bishop [2].

In arriving at a ‘decision’ - *halting* - standard implementations of SDS examine the *stability* of the agent population as a whole; in this manner halting is defined as a *global* property of the agent population. However such global mechanisms are both less biologically/naturally plausible and more complex to implement on parallel computational systems, than local decision making mechanisms. This paper examines the local quorum sensing behaviour observed in some natural (ant) systems and uses this as the inspiration for two new local termination mechanisms - one mechanism, ‘independent termination’, seeks to implement an asynchronous protocol in SDS that in many ways is quite close to the quorum sensing method used by real ants; a second method - confirmation termination - aims to implement a synchronous local termination protocol with more similarity to the conventional SDS architecture.

2 Stochastic Diffusion Search

SDS is based on distributed computation, in which the operations of simple computational units, or agents, are inherently probabilistic. Agents collectively construct the solution by performing independent searches followed by diffusion/communication of information through the population [28]. Positive feedback promotes better solutions by allocating to them more agents for their exploration. Limited resources induce strong competition from which the largest population of agents corresponding to the best-fit solution rapidly emerges.

In many search problems the solution can be thought of as being composed of many subparts and, in contrast to most Swarm Intelligence methods, SDS explicitly utilises such decomposition to increase the search efficiency of individual agents. Thus in SDS each agent poses a hypothesis about the possible solution and evaluates it *partially* [23]. Successful agents repeatedly test their hypothesis while recruiting unsuccessful agents by direct communication. This creates a positive feedback mechanism ensuring rapid convergence of agents onto promising solutions in the space of all solutions. Regions of the solution space labelled

by the presence of agent clusters *with the same hypothesis* can be interpreted as good candidate solutions. A global solution is thus constructed from the interaction of many simple, locally operating agents forming the *largest cluster of agents with the same hypothesis*. Such a *cluster* is dynamic in nature, yet stable, analogous to, “*a forest whose contours do not change but whose individual trees do*”, [4] [9] [29].

Algorithm 1 Classical - inactive recruitment - SDS

```

1: procedure step(swarm, search_space)
2:   for each agent in swarm do                                     ▷ Diffuse Phase
3:     if not agent.active then
4:       polled_agent = swarm.random_agent()
5:       if polled_agent.active then
6:         agent.hypothesis = polled_agent.hypothesis
7:       else
8:         agent.hypothesis = search_space.random_hypothesis()
9:   for each agent in swarm do                                     ▷ Test Phase
10:    test_result = perform_random_test(hypothesis)
11:    agent.active = test_result

```

Central to the power of SDS (see **Algorithm 1**) is its ability to escape local minima. This is achieved by the probabilistic outcome of the partial hypothesis evaluation in combination with reallocation of resources (agents) via stochastic recruitment mechanisms. Partial hypothesis evaluation allows an agent to quickly form its opinion on the quality of the investigated solution without exhaustive testing [23].

The termination of SDS has historically been defined as a function of the *stability* of the population size of a group of *active* agents. Such methods are termed *global* halting criteria as they are a function of the number of active agents within the total population of agents. Two well documented global methods for determining when SDS should halt are the *Weak Halting Criterion* and *Strong Halting Criterion* [26] [25]; the former is simply a function of the *total* number of active agents and the latter the *total* number of active agents maintaining the *same* hypothesis.

3 Collective decision-making in house hunting ants

A model system for collective decision-making is provided by the process of house-hunting in social insects, such as cavity-nesting ants. These ants cannot modify their nest-site and instead relocate the entire colony if the need arises. The processes by which cavity-nesting ants of the genus *Temnothorax* choose a new nest site and emigrate to it has been well-explored, both empirically and theoretically, and is used as a key model of animal collective decision-making.

The ant emigration process can be summarised thus: If the home nest cavity is damaged or degraded, scouts search for new nest sites. Scouts assess available nests across a number of metrics using a weighted additive strategy [16]. If a nest is judged as unsuitable, a scout continues searching; if a nest is assessed as suitable by a scouting ant (Scout A), this scout will return towards the home nest and recruit a second scout (Scout B) [37]. Scout A will lead Scout B to the new nest by tandem-running. Scout B will then make an independent assessment of the nest, and will either reject it and keep searching, or accept the nest and spend some time in it, before returning home and recruiting a further scout. By this positive feedback process, a good quality nest will accumulate ants [34], [36]. Different ants appear to have differing thresholds for starting recruitment to a nest; this means that even low quality nests can attract some ants, but scouts will accumulate more quickly and to a higher level at higher quality nests [39] [35]. This assessment and recruitment process is terminated when scouts sense that a nest site has reached quorum. Scouts then move into a ‘post-quorum’ behavioural state [32]: they stop leading other scouts by tandem-running, and are no longer willing to be recruited by tandem-running themselves. Instead, they transport brood, queen, other workers to the chosen site. Transported workers do not learn the route between the home nest and the new nest, so are unable to return home, and thus cannot challenge the decision that has been implemented [32]. Transported ants therefore contribute strongly to the quorum by staying in the new nest, so once a few scouts have entered a ‘post-quorum’ state and started transporting, others quickly follow suit.

4 Quorum sensing in house-hunting ants

Quorum sensing is widespread throughout biological systems. When a collective decision is required, a quick and effective way of moving from an information-gathering phase to an implementation phase is to use a quorum threshold. A quorum response can be said to occur when an individual’s probability of exhibiting a behaviour (e.g. choosing a given option) is a sharply nonlinear function of the number of other individuals already performing this behaviour (or having chosen that option) [40]. For house-hunting ants, quorum sensing is central to the decision-making process, as it marks the transition from assessment to implementation. Terminating information-gathering promotes cohesion, which is important for ant colonies that only have one reproductively active queen. For cavity-nesting ants, scouts sense quorum by spending 1-2 minutes in a nest assessing the number of workers present via encounter rate, rather using than indirect cues such as pheromone concentration [34] [31]. Quorum threshold as a proportion of colony size is remarkably constant across a range of colony sizes (c3.5%) [13], and this is intriguing, because the relationship between colony size and cavity size is not simple positive correlation: although larger colonies do inhabit larger cavities in the wild, in laboratory tests both small and large colonies prefer larger cavities, presumably to allow for growth [33], [11] & [22].

Quorum sensing is a separate process from quality assessment and recruitment. This means that the quorum sensing process in effect detects an average quality assessment across many scouts, and has the potential to smooth out differences in individual nest acceptance thresholds [17]. Once quorum is reached, scouts do not re-assess quorum on subsequent visits - they will continue to bring brood even if the nest is artificially emptied of ants [31]. The quorum threshold itself is not modulated depending on the quality of the new nest [30]. If nest quality is artificially manipulated during the assessment phase of an emigration, the ants are able to respond flexibly to the new nest qualities; if quality is manipulated after quorum is reached and implementation has begun, then colonies often become ‘trapped’ in an inferior nest [15], [?]. This indicates that quality is not re-assessed after quorum has been reached in these cavity-dwelling *Temnothorax* species. In contrast, a different ant species, *Diacamma indicum* recruits only by tandem-running with no clear quorum point and no adult transport [3]. Colonies of *Diacamma indicum* are able to respond flexibly to manipulated qualities at any stage of the emigration - but overall colony cohesion is lower, supporting the idea that using a quorum threshold increases cohesion, but at a cost to flexibility.

Although quorum sensing behaviour is not modulated by the quality of the options available, it is influenced by the experience and context. Naive scouts use different quorum thresholds to those used by more experienced scouts, but the direction of this difference differs between species, indicating a learning component to quorum sensing behaviour [30]. Emigrations often occur in an emergency context, but cavity-nesting ants do also sometimes emigrate even when their home nest is undamaged, if a better nest is available in the neighbouring area. This is not due to direct comparison of the quality of the two nests, but due to quality-dependent nest acceptance [36], [39]. In these non-emergency migrations, scouts appear to use a quorum threshold around twice as high as in emergency migrations [12], suggesting that colonies prioritise speed over accuracy when conditions are harsher.

5 SDS Local Halting criteria

Drawing inspiration from the behaviour of *Temnothorax* ants in their nest selection behaviour, the halting behaviour of SDS was modified such that it would emerge from purely *local* interactions of SDS agents. By analogy with the behaviour of tandem running *Temnothorax* ants (as outlined in Section (3) and Section (4) herein), in the following we propose two new variants of the process for determining when an agent should switch from the classical SDS *explore-exploit* behaviour to a new, so called, ‘*terminating*’ behaviour which we term the *independent* and *confirmation* halting criteria.

In these variants agents can take on an additional behaviour in which they enter a new state we define ‘*Terminating*’, wherein their hypothesis becomes fixed and they subsequently seek to actively remove agents from the dynamic

swarm³ and give them their own (now fixed) termination hypothesis (analogous to the ‘post-quorum’ behavioural state in ants of the genus *Temnothorax*, wherein post-quorum ants literally carry other ants they encounter to the new nest site).

As this decision making process successively removes agents from the population we name this form of SDS *Reducing SDS*; in this vein a collective ‘decision’ is made (and the local halting condition met) when all agents are either active and/or have been removed from the population.

5.1 ‘Independent’ termination behaviour

In *independent reducing SDS* we relax the assumption that all SDS agents update synchronously in iterative ‘cycles’ (wherein one such cycle corresponds to all agents being updated).

Algorithm 2 Independent SDS

```

1: procedure step(swarm, search_space)
2:   swarm = shuffle(swarm)
3:   for each agent in swarm do
4:     polled_agent = swarm.random_agent()           ▷ Diffusion behaviour
5:     if Both agents are inactive then
6:       Both agents randomise hypothesis
7:     else if One agent is inactive and other is active but not terminating then
8:       Inactive agent assumes active agent’s hypothesis
9:     else if One agent is terminating then
10:      Other agent is removed from the swarm
11:     else if Agents share a hypothesis then
12:       Both agents become terminating
13:     if not agent.terminating then                 ▷ Testing behaviour
14:       test_result = perform_random_test(hypothesis)
15:       agent.active = test_result

```

In independent SDS agents update independently and probabilistically⁴ - which is more analogous to the behaviour of a collection of real ants - and recruitment becomes bidirectional. Considering two such interacting agents:

- if neither agent is active both reselect new random hypotheses;

³ Standard SDS has previously been shown to be a global search algorithm [26] - it will eventually converge to the global best solution in a given search space; by removing agents from the swarm, relative to standard SDS the number of potential agents remaining available for explore-exploit behaviour is reduced; precisely how this reduction impacts the robustness of the algorithm [with respect to erroneous convergence to sub-optimal solutions] has yet to be fully established.

⁴ To facilitate the use of homogenous performance metrics, we assume that in a population of k agents, k single asynchronous updates corresponds to one standard synchronous iteration cycle.

- one agent is inactive and other is active but not terminating then the inactive agent assumes active agent’s hypothesis; in *Temnothorax* nest selection this is analogous to a ‘scout’ ant being recruited to a new nest hypothesis by tandem running.
- one of the agents is in *terminating* mode then the other is assigned the solution hypothesis and ‘removed’ from the population (playing no further part in the search); in *Temnothorax* nest selection this is analogous to an ant in post-quorum (terminating) mode carrying an ant to the new nest site.
- if the two agents meet that both have the same hypothesis then both switch to *terminating* mode; in *Temnothorax* nest selection this recruitment behaviour would serve to reinforce an ant’s initial nest-judgement⁵.

The above process is algorithmically outlined in **Algorithm 2**.

5.2 ‘Confirmation’ termination behaviour

Since its inception in 1989 [7] a substantial body of algorithmic analysis (describing the theoretical behaviour of SDS), empirical studies and practical applications have been published (for a recent review see [2]). To more readily facilitate the future use of these results in both local termination variants and potentially to extend the reach of this analysis to some aspects of real ant behaviour, we suggest a further simplification of Independent SDS to a second reducing behaviour that more closely aligns with standard SDS diffusion; we term this mode *confirmation* reducing termination.

In ‘confirmation reduction’ SDS agents are once again assumed to update *synchronously* and the diffusion of information is changed to more closely resemble that of classical dual mode (passive and active) recruitment SDS [24]. In *confirmation* SDS an active agent polls random agents in the diffusion phase. Active agents become *terminating* if their polled agent is also active and both agents share a hypothesis. The agent is then locked into being active, maintaining that hypothesis. If an inactive agent polls a terminating agent, the inactive agent is removed from the population (see **Algorithm 3** for details).

6 Experiments

A series of experiments was performed to investigate the diffusion behaviour of the two new halting criteria over a variety of search parameters to establish (a) if the algorithms’ gross behaviour remains characteristic of SDS and (b) to

⁵ *Temnothorax* ants are indeed sometimes recruited back to nests they have already visited, so there is potential for this ‘reinforcement recruitment’ process to play a role for ant colonies. For example, ‘reinforcement recruitment’ could cause ants to enter a post-quorum state at a lowered encounter rate. This would help extra rapid acceptance of a nest if there were only one new nest site available. This idea could be tested empirically, ideally in a complex arena that would promote tandem-running behaviour, allowing communication of preference.

Algorithm 3 Confirmation SDS

```

1: procedure step(swarm, search_space)
2:   for each agent in swarm do ▷ Diffuse Phase
3:     polled_agent = swarm.random_agent()
4:     if agent.active then
5:       hyp_1 = agent.hypothesis
6:       hyp_2 = polled_agent.hypothesis
7:       if polled_agent.active and hyp_1 == hyp_2 then
8:         agent.terminating == True
9:       else
10:        if polled_agent.active then
11:          if polled_agent.terminating then
12:            swarm.remove(agent)
13:          else
14:            agent.hypothesis = polled_agent.hypothesis
15:          else
16:            agent.hypothesis = search_space.random_hypothesis()
17:   for each agent in swarm do ▷ Test Phase
18:     test_result = perform_random_test(hypothesis)
19:     agent.active = test_result

```

evaluate their robustness over a variety of search parameters (which effectively characterise the quality of the putative best solution, α ($0 \leq \alpha < 1$), relative to β , ($0 \leq \beta < 1$), the quality of the distractor solution⁶); in the ‘ant migration’ problem, α is analogous to a measure of the quality of the potential new nest site and β effectively a measure of the quality of the original nest.

In all experiments the population is initialised with one agent maintaining the hypothesis representing the potential best solution and the probability of an agent randomly selecting the hypothesis of the potential best solution is set to zero; this ensures that only the *diffusion* behaviour of the algorithm is explored⁷.

In the first experiment each of the three termination functions (strong, independent and confirmation) was modelled in a population of 10000 agents, one of which was active and at the solution hypothesis at time zero, with all other agents set inactive pointing to the ‘noise’ hypothesis. The algorithm was then evaluated 25 times from these conditions against a range of possible values of α and β (from 0 to 0.875 with a step of 0.125). The number of times the algorithm successfully halted within 250 iterations was recorded as was the mean average number of iterations before halting in these cases.

In the case of strong halting SDS, halting was considered successful if the halting criterion was satisfied; figure 1 shows the characteristic *S-shape* con-

⁶ β defines a “uniform random noise” hypothesis; an aggregate of all the possible hypotheses an agent could have other than the putative solution hypothesis.

⁷ These parameters define a problem analogous to the search space being infinitely large, wherein the only way an agent can adopt the ‘best’ solution is to receive it via diffusion from an active agent.

vergence curve obtained deploying SDS using the Strong Halting criterion. All algorithms would also halt if all agents were active at the solution hypothesis, as this is analogous to a successful migration of agents to an optimal state; in addition the process was also halted if the algorithm had run for more than a specified number of iterations or if all the agents held the noise hypothesis. Any experiment that halted for the latter two reasons was considered **unsuccessful**.

In a second experiment the three algorithms were run against fixed values of α and β which the first experiment had shown would be likely to successfully halt. The state of all agents was recorded at every iteration and number of agents (as a proportion of the total population) in various states was graphed over time to visualise the characteristic behaviour of the halting criteria (see figure 2).

Table 1 lists i the average number of iterations before halting and c the number of times that the algorithm successfully halted for a SDS experiment for all three algorithms using a population of 10,000 agents across a variety of parameter values of the noise hypothesis (β) and solution hypothesis (α).

NB. Pairs of values for α and β for which *all three* algorithms failed to converge 25 times out of 25 are not listed. Examining the results presented in **Table 1**, the following comparative observations can be made:

strong halting versus independent reduction on average the convergence time is 3.3 times faster for independent whilst its robustness is similar (strong halting is more robust in 11 cases, less robust in 6);

strong halting versus confirmation reduction on average the convergence time is around 2.8 times faster for confirmation whilst its robustness is similar (strong halting is more robust in 8 cases, less robust in 6);


7 Conclusion

This paper has looked at cooperative decision making in the Stochastic Diffusion Search algorithm and house-hunting ants. Typically, decision making in standard SDS is based on the use of a global halting function which entails *global* access to the activity of the SDS agent population as a whole. Conversely in this paper, inspired by the quorum sensing mechanism deployed by some species of ants in nest selection, we have successfully demonstrated two new *local* termination criteria for SDS which, in terms of their robustness to noise, have been demonstrated to have broadly similar behaviour to the standard SDS meta-heuristic. Furthermore, it is observed that the use of a local halting mechanism results in an approximately three-fold speed-up in the collective decision making time.

Although the independent and confirmation termination processes described in this paper found inspiration from the nest hunting behaviour of *Temnothorax* ants, we do not claim that the nest selection behaviour of these ants is isomorphic to SDS :- one critical difference between the two systems is that SDS relies on its agents being easily able to communicate their current hypothesis to each other, whereas *Temnothorax* ants are only able to do this by the slow [and relatively infrequent] process of tandem-running. Empirical observations have shown that

Table 1. Mean average iterations before termination for three different halting criteria (*strong, independent and confirmation*) over varying quality of solutions

		strong		independent		confirmation	
β	α	i	c	i	c	i	c
0.000	0.625	151.2	17	42.0	6	46.2	14
0.000	0.750	126.8	20	25.8	16	27.8	18
0.000	0.875	118.2	21	20.0	22	21.1	21
0.125	0.625	195.0	7	52.2	12	58.4	7
0.125	0.750	130.4	11	29.6	15	34.1	16
0.125	0.875	122.0	23	22.9	21	25.8	20
0.250	0.625	216.0	1	77.1	7	88.6	5
0.250	0.750	138.7	17	35.5	16	42.2	15
0.250	0.875	125.5	22	26.0	21	31.4	22
0.375	0.625	100.0	1	232.0	1	244.0	1
0.375	0.750	165.6	14	48.5	12	56.7	16
0.375	0.875	131.5	21	30.0	19	38.5	20
0.500	0.750	212.0	5	74.2	12	87.9	7
0.500	0.875	140.5	18	38.5	13	51.5	16
0.625	0.750	100.0	4	150.0	3	238.0	1
0.625	0.875	161.9	18	50.6	14	73.5	20
0.750	0.875	211.0	7	92.4	14	142.5	13
0.875	0.875	100.0	1	—	0	—	0

**Fig. 1.** Cluster size evolution over time for SDS using the *strong halting criterion*. The x-axis counts iterations, the y-axis shows cluster size as a proportion of the entire population. The positive feedback effect can be seen in the sharp *S*-curve of the solution cluster size.




Fig. 2. Cluster size evolution over time for SDS using *confirmation SDS* (left) and *independent SDS* (right). The x-axis counts iterations, the y-axis shows cluster size as a proportion of the entire population. Both show an accelerating growth in the number of agents at the solution hypothesis followed by a similar growth of terminating agents at the solution hypothesis until the entire swarm is either active (in the case of confirmation SDS) or removed from the swarm (in the case of independent SDS).

scouting ants can judge quorum to have been reached (through encounter rate) without having followed a tandem run, so clearly *Temnothorax* ants do not solely rely on *independent-SDS* like termination rules. In this context, future research will investigate the degree to which appropriately modified SDS characterisations can be used to describe the behaviour of *Temnothorax* ants.

References

1. Aleksander, I. & Stonham, T.J.: Guide to pattern recognition using random access memories. *Computers & Digital Techniques* 2(1), pp. 29-40 (1979)
2. al-Rifaie, M.M. & Bishop, J.M.: Stochastic Diffusion Search Review, *Journal of Behavioral Robotics* 4(3), pp. 155-173 (2013)
3. Anoop, K. & Sumana, A.: Response to a change in the target nest during ant relocation. *The Journal of Experimental Biology* 218(6), pp. 887-92 (2015)
4. Arthur, W.B.: Inductive reasoning and bounded rationality, (The El Farol Problem). *Amer. Econ. Rev.* 84, pp. 406-411 (1994)
5. Back, T.: *Evolutionary Algorithms in Theory and Practice*. Oxford University Press, Oxford (1996)
6. Beattie, P.D. & Bishop, J.M.: Self-localisation in the ‘SENARIO’ Autonomous Wheelchair. *Journal of Intelligent and Robotic Systems* 22, pp. 255-267 (1998)
7. Bishop, J.M.: Stochastic Searching Networks. In: *Proc 1st IEE Int. Conf. Artificial Neural Networks*. IEE Conference Publication (313), IEE, London, pp. 329-331 (1989)

8. Bishop, J.M. & Torr, P.H.S.: The Stochastic Search Network. In: Linggard R., Myers, D.J. & Nightingale, C. (eds) *Neural Networks for Images, Speech and Natural Language*. Chapman Hall, New York (1992)
9. Bishop, J.M., Nasuto, S.J. & De Meyer, K.: Dynamic Knowledge Representation in Connectionist Systems. In: Dorrnsoro, J.R. (ed) *Artificial Neural Networks ICANN*, Madrid, 2002. *Lecture Notes in Computer Science*, vol 2415, pp. 308-313. Springer, Heidelberg New York (2002)
10. Bonabeau, E., Dorigo, M. & Theraulaz, G.: *Swarm Intelligence: from Natural to Artificial Systems*. Oxford University Press, Oxford (1999)
11. Cao, T.T.: High social density increases foraging and scouting rates and induces polydomy in *Temnothorax* ants. *Behavioral Ecology and Sociobiology* 67(11), pp. 1799-1807 (2013)
12. Dornhaus, A., Franks, N.R., Hawkins, R.M. & Shere, H.N.S.: Ants move to improve: colonies of *Leptothorax albipennis* emigrate whenever they find a superior nest site. *Animal Behaviour* 67(5), pp. 959-963 (2004)
13. Dornhaus, A., Holley, J.A., Pook, V. G., Worswick, G. & Franks, N.R.: Why do not all workers work? Colony size and workload during emigrations in the ant *Temnothorax albipennis*. *Behavioral Ecology and Sociobiology* 63(1), pp. 43-51 (2008)
14. Franklin, E. L.: The journey of tandem running: the twists, turns and what we have learned. *Insectes Sociaux* 61, pp. 1-8 (2014)
15. Franks, N.R., Hooper, J.W., Gumm, M., Bridger, T.H., Marshall, J.A.R., Gro, R. & Dornhaus, A.: Moving targets: collective decisions and flexible choices in house-hunting ants. *Swarm Intelligence* 1(2), pp. 81-94 (2007).
16. Franks, N.R., Mallon, E.B., Bray, H.E., Hamilton, M.J. & Mischler, T.C.: Strategies for choosing between alternatives with different attributes: exemplified by house-hunting ants. *Animal Behaviour* 65, pp. 215-223 (2003)
17. Franks, N.R., Stuttard, J.P., Doran, C., Esposito, J.C., Master, M.C., Sendova-Franks, A.B., Masuda, N. & Britton, N.F.: How ants use quorum sensing to estimate the average quality of a fluctuating resource. *Scientific Reports* 5: 11890 (2015)
18. Goldberg, D.: *Genetic Algorithms in search, optimization and machine learning*. Addison Wesley, Reading MA (1989)
19. Grech-Cini, E.: Locating facial features. PhD Dissertation, University of Reading, Reading UK (1995)
20. Holland, J.H.: *Adaptation in natural and artificial systems*. The University of Michigan Press, Ann Arbor (1975)
21. Kennedy J, Eberhart R.C. & Shi, Y.: *Swarm Intelligence*. Morgan Kauffman, San Francisco (2001)
22. Kramer, B.H., Scharf, I. & Foitzik, S.: The role of per-capita productivity in the evolution of small colony sizes in ants. *Behavioral Ecology and Sociobiology* 68(1), pp. 41-53 (2013)
23. De Meyer, K., Nasuto, S.J. & Bishop, J.M.: Stochastic Diffusion Optimisation: the application of partial function evaluation and stochastic recruitment. In: Abraham A, Grosam C, Ramos V (eds) *Stigmergic Optimization, Studies in Computational Intelligence* 31, pp. 185-207, Springer, Berlin Heidelberg New York (2006)
24. Myatt, D, M., Nasuto, S.J., Bishop J.M.: Alternative recruitment strategies for SDS. In *Proc. AISB06: Symposium on Exploration vs. Exploitation in Naturally Inspired Search*, pp. 181-187, Bristol UK (2006)
25. Nasuto, S.J.: *Analysis of Resource Allocation of Stochastic Diffusion Search*. PhD Dissertation, University of Reading, Reading UK (1999)
26. Nasuto, S.J. & Bishop, J.M.: Convergence of the Stochastic Diffusion Search. *Parallel Algorithms and Applications* 14, pp. 89-107 (1999)

27. Nasuto, S.J., Bishop, J.M. & Lauria, S.: Time Complexity of Stochastic Diffusion Search. In: Heiss, M. (ed) Proc. Int. ICSC / IFAC Symposium on Neural Computation. Vienna (1998)
28. Nasuto, S.J., Dautenhahn, K. & Bishop, J.M.: Communication as an Emergent Metaphor for Neuronal Operation. In: Nehaniv, C. (ed) Computation for Metaphors, Analogy, and Agents, Lecture Notes in Artificial Intelligence 1562, pp. 365-379, Springer, Berlin Heidelberg New York (1999)
29. Nasuto, S.J., Bishop, J.M. & De Meyer, K.: Communicating neurons: a connectionist spiking neuron implementation of stochastic diffusion search, *Neurocomputing* 72(4-6), pp. 704-712, Elsevier, Orlando (2008)
30. Pratt, S.C.: Behavioral mechanisms of collective nest-site choice by the ant *Temnothorax curvispinosus*. *Insectes Sociaux* 52, pp. 383-392 (2005)
31. Pratt, S.C.: Quorum sensing by encounter rates in the ant *Temnothorax albigipennis*. *Behavioral Ecology* 16, pp. 488-496 (2005)
32. Pratt, S.C., Mallon, E.B., Sumpter, D.J.T. & Franks, N.R.: Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant *Leptothorax albigipennis*. *Behavioral Ecology and Sociobiology* 52(2), pp. 117-127 (2002)
33. Pratt, S.C. & Pierce, N.E.: The cavity-dwelling ant *Leptothorax curvispinosus* uses nest geometry to discriminate among potential homes. *Animal Behaviour* 62: 281-287 (2001)
34. Pratt, S.C., Sumpter, D.J.T., Mallon, E.B. & Franks, N.R.: An agent-based model of collective nest site choice by the ant *Temnothorax albigipennis*. *Animal Behaviour* 70, pp. 1023-1036 (2005)
35. Robinson, E.J.H., Feinerman, O. & Franks, N.R.: How collective comparisons emerge without individual comparisons of the options. *Proc. Royal Soc. B* 2014 281 20140737; DOI: 10.1098/rspb.2014.0737. (2014)
36. Robinson, E.J.H., Franks, N.R., Ellis, S., Okuda, S. & Marshall, J.A.R.: A simple threshold rule is sufficient to explain sophisticated collective decision-making. *PLoS One* 6: e19981. (2011)
37. Robinson, E.J.H., Smith, F.D., Sullivan, K.M.E & Franks, N.R.: Do ants make direct comparisons? *Proc. Royal Soc. B* 276: 2635-2641. (2009)
38. Seeley, T.D., Visscher, P.K., Schlegel, T., Hogan, P.M., Franks, N.R. & Marshall J.A.R.: Stop signals provide cross inhibition in collective decision-making by honey bee swarms. *Science* 335, pp. 108-111 (2012)
39. Stroeymeyt, N., Robinson, E.J.H., Hogan, P.M., Marshall, J.A.R., Giurfa, M. & Franks, N.R.: Experience-dependent flexibility in collective decision-making by house-hunting ants. *Behavioral Ecology* 22(3), pp. 535-542 (2011)
40. Sumpter, D.J.T. & Pratt, S.C.: Quorum responses and consensus decision making. *Proc. Royal Society B* 364(1518), pp. 743-753 (2009)
41. Whitaker, R.M. & Hurley, S.: An agent based approach to site selection for wireless networks. In: Proc. 2002 ACM Symp. Applied Computing (Madrid), pp. 574 - 577, ACM, New York (2002)