

The mining game: a brief introduction to the Stochastic Diffusion Search metaheuristic*

In recent years studies of social agents have suggested several new metaheuristics for use in search and optimisation; Stochastic Diffusion Search (SDS) [1] is one such ‘Swarm Intelligence’ algorithm. SDS is a distributed population based search algorithm utilising interaction between simple agents to locate a global optimum; such ‘communicating agents’ have recently been suggested as a potential metaphor for some cognitive processes [6].

SDS is most easily applied to discrete search and optimisation problems where the task is to identify the hypothesis, h , which maximises the value of a *decomposable objective function*¹. Unlike many *nature inspired search algorithms* SDS has a solid mathematical framework which fully describes the behaviour of the algorithm, investigating its: resource allocation [4], convergence to global minima [5], robustness and minimal convergence criteria [2] and time complexity [3]. In the following brief summary we deploy a simple new metaphor - *the mining game* - to introduce SDS to readers of the AISBQ.

The mining game provides a high-level description of a search to identify the best hill, H_{best} , in a large mountain range on which a group of miners should prospect for gold; each hill is quantised into a fixed set of regions, R , where each region yields a specific ‘rate of return’ R_j of gold (*concentration*). Thus the ‘best’ hill for the miners is the hill, H_i , which maximises the value of the [decomposable] objective function $F = H_i \sum_j R_j$.

The mining game: a group of miners learn that there is gold to be found on the hills of a large mountain range but have no information regarding its distribution. To maximize their collective wealth, the maximum number of miners should dig where the concentration of gold is highest; but this information is not available a-priori. Thus the goal of the mining game [*resource allocation process*] is to allocate the most miners to the hill which has the richest seams of gold. In order to solve this problem, the miners employ a simple Stochastic Diffusion Search.

*Mohammad Majid al-Rifaie and Mark Bishop, Department of Computing, Goldsmiths College, University of London. London SE14 6NW, UK.

¹ A *decomposable objective function* F is one which can be independently evaluated via a set of partial functions, f_i , such that $F = \sum_i f_i$.

At the start of the mining process each miner is randomly allocated a hill to mine (his hypothesis, h). Every day each miner digs a randomly chosen region on the hill, R_j . At the end of each day the probability that a miner is ‘successful’ is proportional to the amount of gold he has mined (the **test** phase). Each evening the miners congregate and each miner who is not successful selects another miner at random for communication. If the chosen miner has been successful then they share the location of the gold and subsequently both maintain it as their hypothesis, h ; if not, the first [unsuccessful] miner selects a new region to mine at random (the **diffusion** phase).

By iterating through test and diffusion phases miners stochastically explore the whole solution space. However, since tests succeed more often on rich seams than on poor, an individual miner will spend more time examining *good* regions, at the same time recruiting other miners, which in turn recruit even more miners. Candidate solutions are thus identified by concentrations of a substantial population of miners.

Central to the power of SDS 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 - *miners* - via stochastic recruitment mechanisms. Partial hypothesis evaluation allows a miner to quickly formulate an ‘opinion’ on the quality of the investigated solution without exhaustive testing (e.g. a miner forms an opinion on the best hill in the range without exhaustively digging for gold in each region on every hill).

At each iteration of the algorithm the ‘success’ of the miners in their search for gold can be evaluated probabilistically or deterministically; in the former case each region has a probability of yielding gold, in the latter gold is either present or absent at each region in discrete parcels. In both cases at the end of the test phase miners are either successful or unsuccessful (cf. standard SDS); thus the mining game can be further refined through either of the following two assumptions:

1. *Finite resources*: the probability of finding gold is reduced each time a miner successfully mines a region;
2. *Infinite resources*: the probability of finding gold in a region does not vary as gold is extracted from it.

In the first case - finite resources - the task is dynamic and analagous to the robot search tasks described by Steels [8] and Krieger [7]; the second case is analogous to conventional discrete optimisation problems. Ongoing research seeks to apply SDS to the wider gamut of continuous optimisations problems - revisiting the Mining Game metaphor, this is analagous to locating the [real-valued] point(s) on the mountain range where the concentration of gold is highest. In Cognitive Science other recent work offers ‘stochastic diffusion’ as a new potential metaphor for neuronal operation [6].

References

- [1] J. M. Bishop, Stochastic Searching Networks, Proc. 1st IEE Int. Conf. ANNs, London UK, pp. 329-331, 1989.
- [2] D. M. Myatt, J. M. Bishop and S. J. Nasuto, Minimum stable convergence criteria for Stochastic Diffusion Search, Electronics Letters 22(40), pp. 112-113, 2004.
- [3] S.J. Nasuto, J.M. Bishop & S. Lauria, Time Complexity of Stochastic Diffusion Search. Neural Computation '98, pp. 260-266, Vienna Austria, 1998.
- [4] S. J. Nasuto, Resource Allocation Analysis of the Stochastic Diffusion Search, PhD Thesis, The University of Reading, Reading UK, 1999.
- [5] S.J. Nasuto and J. M. Bishop, Convergence Analysis of Stochastic Diffusion Search, Journal of Parallel Algorithms and Applications 14(2), pp: 89-107, 1999.
- [6] S.J. Nasuto, J. M. Bishop, and K. De Meyer, Communicating neurons: a connectionist spiking neuron implementation of stochastic diffusion search, Neurocomputing (72), pp. 704-712, 2009.
- [7] M. J. Krieger, J. B. Billeter, and L. Keller, Ant-like task allocation and recruitment in cooperative robots, Nature, vol. 406 (6799), pp. 9925, 2000.
- [8] L. Steels, Cooperation between distributed agents through selforganization, Proc. IEEE Int Workshop on Intelligent Robots and System (7), pp. 814, 1990.