

Swarmic Paintings and Colour Attention

Mohammad Majid al-Rifaie¹ and John Mark Bishop²

¹ Vividus Solutions LTD., Vividus Swarm Lab, London EC1V 2NX
m.majid@vividusolutions.com

² Goldsmiths, University of London, New Cross, London SE14 6NW
m.bishop@gold.ac.uk

Abstract. Swarm-based multi-agent systems have been deployed in non-photorealistic rendering for many years. This paper introduces a novel approach in adapting a swarm intelligence algorithm – Stochastic Diffusion Search – for producing non-photorealistic images. The swarm-based system is presented with a digital image and the agents move throughout the digital canvas in an attempt to satisfy the dynamic roles – attention to different colours – associated to them via their fitness function. Having associated the rendering process with the concepts of ‘attention’ in general and colour attention in particular, this papers briefly discusses the ‘computational creativity’ of the work through two prerequisites of creativity (i.e. freedom and constraints) within the swarm intelligence’s two infamous phases of exploration and exploitation.

1 Introduction

In recent years, studies of the behaviour of social insects (e.g. ants and bees) and social animals (e.g. birds and fish) have proposed several new metaheuristics for use in collective intelligence. Natural examples of swarm intelligence that exhibit a form of social interaction are fish schooling, birds flocking, ant colonies in nesting and foraging, bacterial growth, animal herding, brood sorting etc.

Although producing artistic works through the use of swarm intelligence techniques have been previously explored, this work explores the concepts of attention and creativity through this type of collective intelligence, which emerges through the interaction of simple agents, representing the social insects and animals, in a nature-inspired algorithm – Stochastic Diffusion Search (SDS) [8].

The swarm intelligence algorithm adapted for the present work, utilises swarms with dynamically changing ‘attention’, exhibited via iteratively attending-to (cf. (re)painting) different colours in the source image. The agents thus converge their attention on areas with similar colours in the source image (search space). This process is repeated and the outcome – emerging through millions of simple interactions – constantly changes based on how the swarm controls its attention.

Following other works in the field of swarms painting ([7,15,19,20]) and ant colony paintings ([9,14]), the outputs presented in this paper – created by a swarm intelligence algorithm – are used as a platform to argue whether or not swarm intelligence algorithms have the potential to exhibit computational creativity.

In this paper, first the swarm intelligence algorithm used is explained, and subsequently a historical perspective of attention is presented followed by explanation on how colour attention can possibly be used in creating artistic works. Some technical information are detailed afterwards on the performance of the computer-generated nature-inspired colour-attentive swarms in rendering images. Then a short discussion follows on whether swarms can be computationally creative, with reference to freedom and constraint, which are associated to the two well-known phases of exploration and exploitation. The conclusion appears at the end along with some suggestions for possible future research.

2 Stochastic Diffusion Search

This section describes a swarm intelligence algorithm (Stochastic Diffusion Search), which is inspired by a species of ants and uses communication as its main mean to converge to an optimum food location by recruiting individual ants. This algorithm is adapted for rendering images in this work.

The performance of Stochastic Diffusion Search (SDS) [8] is based on simple interaction of agents. This algorithm is inspired by one species of ants, *Leptothorax acervorum*, where a ‘tandem calling’ mechanism (one-to-one communication) is used, the forager ant that finds the food location recruits a single ant upon its return to the nest; therefore the location of the food is physically publicised [13].

The SDS algorithm commences a search or optimisation by initialising its population and then iterating through two phases (see Algorithm 1).

Algorithm 1. SDS Algorithm

```
01: Initialise agents
02: While (stopping condition is not met)
04:   For each agent
03:     Test hypothesis and determine activity
05:   For each agent
06:     Diffuse hypothesis
07: End While
```

In the test phase, SDS checks whether the agent hypothesis is successful or not by performing a hypothesis evaluation which returns a boolean value. Later in the iteration, contingent on the precise recruitment strategy employed (in the diffusion phase), successful hypotheses diffuse across the population and in this way information on potentially good solutions spreads throughout the entire population of agents. In other words, each agent recruits another agent for interaction and potential communication of hypothesis.

In standard SDS (which is used in this paper), *passive recruitment mode* is employed. In this mode, if the agent is inactive, a second agent is randomly selected for diffusion; if the second agent is active, its hypothesis is communicated (*diffused*) to the inactive one. Otherwise there is no flow of information

between agents; instead a completely new hypothesis is generated for the first inactive agent at random (see Algorithm 2). Therefore, recruitment is not the responsibility of the active agents.

Algorithm 2. Passive Recruitment Mode

```

01: For each agent
02:   If ( !ag.isActive )
03:     r_ag = pick a random agent
04:     If ( r_ag.isActive )
05:       ag.hypothesis = r_ag.hypothesis
06:     Else
07:       ag.hypothesis = generate a random hypothesis
08:   End If
09: End For

```

3 Attention

The concept of attention has been studied mostly in psychology and neuroscience (see Table 1.1 in Phuong Vu: Historical Overview of Research on Attention, in [21] for more details) and there has been considerably less notable interest on attention within the field of computational creativity.

In the early 18th century attention was mostly seen as a way of abstraction (see Berkeley’s 1710 theory of abstract ideas in [16]):

“[It] must be acknowledged that a man may consider a figure merely as triangular, without attending to the particular qualities of the angles or relations of the sides. So far he may abstract, but this will never prove that he can frame an abstract general, inconsistent idea of a triangle. ”

By 1769, when Henry Home Kames added the appendix of ‘Terms Defined or Explained’ to his *Elements of Criticism* [12], attention’s role as a regulator of cognitive input was regarded as definitive of it:

“Attention is that state of mind which prepares one to receive impressions. According to the degree of attention objects make a strong or weak impression. Attention is requisite even to the simple act of seeing.”

Thus, regulating cognitive and sensory inputs was associated to attention. Later, William James in *The Principles of Psychology* in 1890 [10] offered a more comprehensive description of attention (i.e. focalisation, etc.):

“Every one knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence [...]” (p. 403-404)

and few pages further, he continues:

“Each of us literally chooses, by his ways of attention to things, what sort of a universe he shall appear to himself to inhabit.” (p. 424)

Two decades later, in 1908, as emphasised by E.B. Titchener [17], attention was given a greater significance:

“What I mean by the ‘discovery’ of attention is the explicit formulation of the problem: the recognition of its separate status and fundamental importance; the realization that the doctrine of attention is the nerve of the whole psychological system, and that as men judge of it, so they shall be judged before the general tribunal of psychology.”

and its importance grew over the years in psychology and neuroscience. Although the concept of attention might have been present in the work of some researchers in the field of computational creativity, the focus on attention has not been equally considerable among researchers in this field; perhaps, partly because there has not been a clear definition on attention.

The next section adopts a particular type of attention (i.e. colour attention) and expands on its application in the context of painting swarms (or *swarmic paintings*).

4 Colour Attention and Creativity in the Swarms

In this section, after describing an artistic technique (painting by numbers), the technical details of the adapted Stochastic Diffusion Search algorithm is given and the results are presented in form of the images rendered by the swarms.

4.1 Painting by Numbers

In this initial research, the authors aim at addressing ‘colour attention’ by utilising a swarm intelligence algorithm whose agents dynamically render an input image, solely by communicating the colour qualities of various pixels of the input image.

The final product of the above mentioned approach could be assimilated to a well-known technique – Painting by Numbers – invented by Max S. Klein in 1950 (kits having a board on which light blue or gray lines indicate areas to paint; each area is labelled with a number which in turn is associated to a colour). One of the most famous artist known for using this technique is Andy Warhol [18].

In this work, and following the concept of the aforementioned technique, colour guides the ‘attention’ of the swarm (i.e. colour is the input of the fitness function). If there is a large region within a painting, it is more likely that it would be noticed by the viewer before the smaller regions. As it will be presented next, the behaviour of the swarm is also influenced by this factor.



Fig. 1. Input Image

4.2 Colour-Attentive Swarms

As mentioned earlier in Section 2, each agent has two components: status, which is a boolean value, and hypothesis. The hypothesis of each agent in this work is the (x, y) coordinate pointing to the colour attributes (i.e. RGB values) of a particular pixel within the input image (search space).

In the beginning of each iteration, a pixel is chosen and its corresponding colour attributes (focal colour or f_c) are kept as the focal attention of the swarm.

In order to determine the status of the agents within the swarm (test phase), the colour distance (d_{i_c}) of each agent's colour from f_c is calculated according to the Eq. 1; if $d_{i_c} < \alpha$, the agent is set to be active, otherwise inactive.

$$d_{i_c} = d(f_c, A_{i_c}) = \sqrt{(R_{f_c} - R_{A_{i_c}})^2 + (G_{f_c} - G_{A_{i_c}})^2 + (B_{f_c} - B_{A_{i_c}})^2} \quad (1)$$

where A_{i_c} is the colour associated to the i^{th} agent; R_{f_c} , G_{f_c} , B_{f_c} and $R_{A_{i_c}}$, $G_{A_{i_c}}$, $B_{A_{i_c}}$ are the RGB values of f_c and the i^{th} agent respectively.

In the diffusion phase, as in standard SDS, each inactive agent randomly pick another one. If the randomly selected agent is active, the inactive agent adopts (a Gaussian random distance, $\sigma = 5$, and direction from) the (x, y) coordinates of the active agent, allowing to explore a small area around the active agent; the colour associated with the active agent is painted on the canvas with varying diameters (see Section 4.3 for details). However, if the selected agent is inactive, the selecting agent generates a random (x, y) coordinates (from the search space), whose status would be defined in the test phase of the next iteration. New f_c is generated when either 10 iterations (test-diffusion cycles) are run or the entire agents population become active.




Fig. 2. Stages through which the attention of the swarm is shifted from one colour to the next (top to bottom); note that the images of each column are rendered independently using the focal colour (f_c , on the leftmost of the figure) and the α values provided (5, 15 and 25)

The value of α determines how focus the attention of the swarms should be on a particular colour (f_c); the greater the value of α , the bigger the colour-similarity ‘tolerance’; and thus more agents residing within the colour range of f_c would be potentially active. Smaller values of α ensure higher attention on the focal colour (f_c) and therefore less colour similarity tolerance.

4.3 Experiments and Results

Fig. 1 shows the original image used as input to the proposed system. The swarms population size is determined by $(w \times h)/5$, where w is the image width and h is the image height. In this work α is set to 5, 15 and 25 in three separate trials; The width and height of the input image is $w = 640$, $h = 428$ respectively and thus, the population size is 54,784.

Once the input image is introduced to the system, f_c is (randomly) generated (one at a time) from the search space (see the colour labels in Fig. 2 – on the left) and then SDS uses the generated f_c to go through the test and diffusion phases, then the next f_c is used and so forth. Fig. 2 (top to bottom) illustrate the visual effect of this process on the digital canvas. Although in this example, f_c values are predefined to allow the three trials ($\alpha = 5, 15$ and 25) use the same f_c values, in practice, the probability of a colour been picked is proportional to its presence in the digital canvas.

As shown, the attention of the swarms is controlled through the generated f_c and value of α . The visual feedback of the swarms on the canvas allows the viewer to observe the change of attention from one colour to the next. The smaller the value of α the more precise (colour-wise) the attention of the swarm. This can be observed in Fig. 2 where three values of α are examined and as the images show, the attention of the swarms is more focused on f_c when α is smaller (e.g. $\alpha = 5$) than when α is larger (e.g. $\alpha = 25$).

In the beginning, the swarms are randomly initialised throughout the search space. The size of the brush is set to be inversely proportional to the activity level of the swarms (i.e. if a high number of the agents become active during the test phase, the brush size is made smaller, and vice versa).

The next part addresses exploration and exploitation as the two main phases in any swarm intelligence algorithms, and expanding on that, a link is made between these phases and the two prerequisites of creativity (i.e. freedom and constraint).

4.4 Freedom vs. Constraint

Both freedom and constraint have always been at the core of several definitions for creativity. Philip Johnson-Laird in his work on freedom and constraint in creativity [11] states:

*“... for to be creative is to be free to choose among alternatives .. []
.. for which is not constrained is not creative.”*

In swarm intelligence systems, the two phases of exploration and exploitation introduce the freedom and control the level of constraint. Pushing the swarms towards exploration, freedom is boosted; and by encouraging exploitation, constraint is more emphasised. Finding a balance between exploration and exploitation has been an important theoretical challenge in swarm intelligence research and over the years many hundreds of different approaches have been deployed by researchers in this field.



Fig. 3. Landscapes. Images on the leftmost column are the input, and the middle and right columns show the output images rendered by the swarms.

In the presented work, one such swarm intelligence algorithm is deployed. This algorithm (i.e. Stochastic Diffusion Search) is responsible for “intelligently” controlling the attention of the agents. This algorithm mimics the behaviour of one species of ants foraging and has an internal mechanism of balancing off the exploitation and exploration phases. Within the diffusion phase of the algorithm, if an inactive agent randomly chooses an active one, the hypothesis of the active agent is diffused to the inactive one (i.e. effectively, the inactive agent is drawn towards the active one). This process boosts exploitation.




Fig. 4. Flame close-ups. Images on the leftmost column are the input, and the middle and right columns show the output images rendered by the swarms.




Fig. 5. Wall texture and flower close-ups. Images on the leftmost column are the input, and the middle and right columns show the output images rendered by the swarms.

On the other hand, if the inactive agent randomly chooses another inactive agent, the selecting agent is randomly restarted within the search space (digital canvas) and thus prompting exploration.

Stochastic Diffusion Search has been used on numerous occasions alongside other swarm intelligence algorithms (e.g. Particle Swarm Optimisation, Differential Evolution Algorithm, etc. in [4,5]). And this paper present the novel approach of deploying *solely* SDS algorithm for the purpose of producing artistic works. As known, there have been several relevant attempts to create creative computer generated artwork using Artificial Intelligence, Artificial Life and Swarm Intelligence. Irrespective of whether the swarms are considered genuinely creative or not, their similar individualistic approach is not totally dissimilar to those of the “elephant artists” [22]:

“After I have handed the loaded paintbrush to [the elephants], they proceed to paint in their own distinctive style, with delicate strokes or broad ones, gently dabbing the bristles on the paper or with a sweeping flourish, vertical lines or arcs and loops, ponderously or rapidly and so on. No two artists have the same style.”

Similarly if the same input image is repeatedly given to the swarms, the images rendered by the swarms at each time, are never the same. In order to evaluate this claim empirically, the performance of the swarms are observed when presented

with several input images. In this experiment, α is set to 15; the value of f_c is randomly generated from the input image and thus the choice of f_c depends on the level of presence of each particular colour in the search space.

Once the process is started, the digital canvas dynamically changes its state and each time produces a different ‘interpretation’ of the input image; this is shown in Figures 3, 4 and 5 where the leftmost images represent the input; on their right, two randomly chosen snapshots of the dynamic work of the swarms rendering their corresponding input images are displayed. While the output images (produced by the freedom and constraints of the swarms) stay loyal to the input image, each is unique and different from the next.

Finally, although this work uses Stochastic Diffusion Search to intelligently (vs. greedily or randomly) control the colour attention of the swarms, the concept of attention is *extendible* to other measures (e.g. shapes); see [3] for an example of using SDS-led attention in producing sketches called ‘Swarmic Sketches’.

5 Conclusion

This paper introduces a novel approach of using Stochastic Diffusion Search (SDS) to generate non-photorealistic images with emphasis on the concept of attention in general, and more specifically colour attention.

The adapted SDS algorithm for the present work, utilises the swarms with dynamically changing colour attention. The swarms thus converge their attention on areas with similar colours in the search space and the attention of the swarms is visualised through their paintings on the digital canvas. This process is repeated and the outcome – emerging through millions of simple interactions – constantly produces (although loyal, yet) non-identical rendering of the same input image.

This work also highlights the mechanisms responsible for the exploration and exploitation phases within the swarm intelligence algorithm and their relationship with freedom and constraints as two prerequisites of creativity. Expanding on the previous research on computational creativity (e.g. [1,2,6]), in addition to the practical aspect of the work presented, the possibility of exhibiting ‘computational creativity by a novel use of the Stochastic Diffusion Search algorithm is briefly discussed.

Acknowledgement. The figures used in this work are the following (in order of appearance):

1. "Port Nouveau," ©2008 David Blaikie
2. "Castello Normanno Bari, Italy," ©2012 L.C. Nøttaasen
3. "gottfried bhm, pilgrimage church, nevigis 1963--1972," ©2008 seier+seier
4. "Untitled," ©2009 Kevin Dooley
5. "Beyond The Flame," ©2012 Tc Morgan
6. "brick wall," ©2009 L S
7. "Untitled," ©2009 Claus Rebler
8. "a bee," ©2009 *n3wjack's world in pixels

Photos no. 1, 2, 3 and 4 are used under a Creative Commons Attribution 2.0 Generic license: <http://creativecommons.org/licenses/by/2.0/>; photo no. 5 is used under a Creative Commons Attribution-NonCommercial-ShareAlike 2.0 Generic license: <http://creativecommons.org/licenses/by-nc-sa/2.0/>; photo no. 6 is used under a Creative Commons Attribution-NonCommercial 2.0 Generic license: <http://creativecommons.org/licenses/by-nc/2.0/>; photos no. 7 and 7 are used under a Creative Commons Attribution-ShareAlike 2.0 Generic license: <http://creativecommons.org/licenses/by-sa/2.0/>.

References

1. Al-Rifaie, M.M., Aber, A., Bishop, M.: Cooperation of nature and physiologically inspired mechanisms in visualisation. In: Ursyn, A. (ed.) *Biologically-Inspired Computing for the Arts: Scientific Data through Graphics*. IGI Global, United States (2012) ISBN13: 9781466609426, ISBN10: 1466609427
2. Al-Rifaie, M.M., Bishop, M.: Weak vs. strong computational creativity. In: *AISB 2012: Computing and Philosophy*, University of Birmingham, Birmingham, U.K. (2012)
3. Al-Rifaie, M.M., Bishop, M.: Swarmic Sketches Deploy Attention Mechanism. In: Machado, P., McDermott, J., Carballal, A. (eds.) *EvoMUSART 2013*. LNCS, vol. 7834, pp. 85–96. Springer, Heidelberg (2013)
4. Al-Rifaie, M.M., Bishop, M., Blackwell, T.: An investigation into the merger of stochastic diffusion search and particle swarm optimisation. In: *GECCO 2011: Proceedings of the 2011 GECCO Conference Companion on Genetic and Evolutionary Computation*, pp. 37–44. ACM (2011)
5. Al-Rifaie, M.M., Bishop, M., Blackwell, T.: Information sharing impact of stochastic diffusion search on differential evolution algorithm. In: *Journal of Memetic Computing*, pp. 1–12. Springer, Heidelberg (2012), <http://dx.doi.org/10.1007/s12293-012-0094-y>, doi:10.1007/s12293-012-0094-y
6. Al-Rifaie, M.M., Bishop, M., Caines, S.: Creativity and Autonomy in Swarm Intelligence Systems. In: Bishop, M., Erden, Y. (eds.) *Journal of Cognitive Computation: Computational Creativity, Intelligence and Autonomy*, vol. 3, pp. 320–331. Springer, Heidelberg (2012)
7. Aupetit, S., Bordeau, V., Monmarche, N., Slimane, M., Venturini, G.: Interactive evolution of ant paintings. In: *The 2003 Congress on Evolutionary Computation, CEC 2003*, vol. 2, pp. 1376–1383 (2004)
8. Bishop, J.: Stochastic searching networks. In: *Proc. 1st IEE Conf. on Artificial Neural Networks*, London, UK, pp. 329–331 (1989)
9. Greenfield, G.: Evolutionary Methods for Ant Colony Paintings. In: Rothlauf, F., Branke, J., Cagnoni, S., Corne, D.W., Drechsler, R., Jin, Y., Machado, P., Marchiori, E., Romero, J., Smith, G.D., Squillero, G. (eds.) *EvoWorkshops 2005*. LNCS, vol. 3449, pp. 478–487. Springer, Heidelberg (2005)
10. James, W.: *The principles of psychology* (1890)
11. Johnson-Laird, P.N.: Freedom and constraint in creativity. In: Sternberg, R.J. (ed.) *The Nature of Creativity: Contemporary Psychological Perspectives*, pp. 202–219. Cambridge University Press (1988)
12. Kames, H.H.: *Elements of Criticism* (1769)

13. Moglich, M., Maschwitz, U., Holldobler, B.: Tandem calling: A new kind of signal in ant communication. *Science* 186(4168), 1046–1047 (1974)
14. Monmarche, N., Aupetit, S., Bordeau, V., Slimane, M., Venturini, G.: Interactive evolution of ant paintings. In: McKay, B., et al. (eds.) 2003 Congress on Evolutionary Computation, vol. 2, pp. 1376–1383. IEEE Press (2003)
15. Moura, L., Ramos, V.: Swarm paintings–nonhuman art. *Architopia Book, Art, Architecture and Science*, 5–24 (2007)
16. Taylor, C.C.W.: Berkeley’s theory of abstract ideas. *The Philosophical Quarterly*, 97–115 (1978)
17. Titchener, E.B.: *Lectures on the elementary psychology of feeling and attention*. The Macmillan Company (1908)
18. Turner, C.: Exhibition: Industrial paint’s artistic legacy. *Nature* 452, 412 (2008)
19. Urbano, P.: Playing in the Pheromone Playground: Experiences in Swarm Painting. In: Rothlauf, F., Branke, J., Cagnoni, S., Corne, D.W., Drechsler, R., Jin, Y., Machado, P., Marchiori, E., Romero, J., Smith, G.D., Squillero, G. (eds.) *EvoWorkshops 2005*. LNCS, vol. 3449, pp. 527–532. Springer, Heidelberg (2005)
20. Urbano, P.: Consensual Paintings. In: Rothlauf, F., Branke, J., Cagnoni, S., Costa, E., Cotta, C., Drechsler, R., Lutton, E., Machado, P., Moore, J.H., Romero, J., Smith, G.D., Squillero, G., Takagi, H. (eds.) *EvoWorkshops 2006*. LNCS, vol. 3907, pp. 622–632. Springer, Heidelberg (2006)
21. Vu, P.: *Historical Overview of Research on Attention*. Sage Publications, Incorporated (2003)
22. Weesatchanam, A.M.: Are Paintings by Elephants Really Art? The Elephant Art Gallery (July 31, 2006)