# A SWARM INTELLIGENCE METHOD FOR FEATURE TRACKING IN AMV DERIVATION

**Angeles Hernandez-Carrascal, Slawomir J. Nasuto**

School of Systems Engineering, University of Reading, Whiteknights, Reading RG6 6AY, U.K.

**Abstract**

Feature tracking is a key step in the derivation of Atmospheric Motion Vectors (AMV). Most operational derivation processes use some template matching technique, such as Euclidean distance or cross-correlation, for the tracking step. As this step is very expensive computationally, often short-range forecasts generated by Numerical Weather Prediction (NWP) systems are used to reduce the search area. Alternatives, such as optical flow methods, have been explored, with the aim of improving the number and quality of the vectors generated and the computational efficiency of the process. This paper will present the research carried out to apply Stochastic Diffusion Search, a generic search technique in the Swarm Intelligence family, to feature tracking in the context of AMV derivation. The method will be described, and we will present initial results, with Euclidean distance as reference.

## 1. INTRODUCTION

Feature tracking is a key step in the derivation of Atmospheric Motion Vectors (AMV), and most operational derivation processes use some template matching technique for this step. Euclidean distance and cross correlation are two widely used template matching techniques, described e.g. in Dew and Holmlund (2000). These methods are computationally expensive, and often a short-range forecast wind field produced by a numerical weather prediction (NWP) system is used to reduce the search area.

Other tracking methods have been explored with the aim of improving the quality and quantity of the generated AMVs and/or decreasing the computational cost of the feature tracking step, which is the most expensive component of the overall AMV derivation process. During the last decade, optical flow methods (e.g. Bresky and Daniels, 2006) have been the main alternative to traditional template matching methods for calculating motion vectors. However, feature tracking has not attracted as much attention as e.g. height assignment from the AMV research community. Although traditionally feature tracking and height assignment have been considered as independent steps in the derivation process, recent research by Borde and Oyama (2008) proposes a link between the two steps, and uses the pixels with the highest contribution to the correlation in feature tracking for the height assignment.

The problem of calculating motion vectors through template matching techniques can be seen as a search problem. Once a target template has been selected in an image i1, we look for its best match in the following image, i2, in a search area determined e.g. by the maximum expected wind or by the maximum expected departure from a forecast wind (if a NWP-generated wind field is used). An objective function, such as Euclidean distance or cross correlation, describes the similarity between the target and potential matches. We can visualise those functions defined on the search space as similarity or correlation landscapes. If cross correlation is used, we look for the global highest hill in the landscape, and if Euclidean distance is used, we look for the global minimum of the landscape.

Stochastic Diffusion Search (SDS), originally developed by Bishop (1989), is a global search technique in the swarm intelligence family. Swarm intelligence techniques are population-based, and the problem-solving abilities of the system emerge from simple individual behaviour and interaction within the collective.

This paper reports on a research project that started recently as a real-life, challenging application of SDS, and whose purpose is to explore the potential of SDS to address feature tracking in the context of AMV derivation. Section 2 describes the SDS algorithm, and chapter 3 discusses SDS as a potential framework to tackle feature tracking. Section 4 presents some early results and section 5 concludes the paper.

## 2. STOCHASTIC DIFFUSSION SEARCH

Stochastic Diffusion Search (SDS) can be applied to search problems for which the objective function that evaluates the quality of candidate solutions can be decomposed into independent elements or micro-features. SDS is based on a population of agents that can perform partial evaluations of candidate solutions to the search problem and can communicate with other agents in a simple way. It was originally developed as a pattern recognition technique to deal with the problem of stimulus-equivalence, and has since then been extended, analysed and applied to real-world problems. A recent description is given by De Meyer et al. (2006). In particular, it has been applied to feature tracking to locate features of human faces in video images (Grech-Cini and McKee, 1993; Bishop and Torr, 1992).

The standard SDS algorithm is shown in pseudo-code in Table 1. Initially, agents are spread randomly over the search space, i.e. each agent selects randomly a candidate solution (or location, or hypothesis) from the search space. After the initialization, there is an iterative process; each of its steps consists of two phases. During the first phase, called test phase, each agent randomly selects a micro-feature, i.e. an element of the objective function, and performs the partial evaluation of the location corresponding to the selected micro-feature; in standard SDS the result of such an evaluation can only be either successful or unsuccessful. If the evaluation is successful the agent becomes active, and otherwise it becomes inactive. Active agents keep their hypotheses for the next iteration. During the second phase, called diffusion of information, inactive agents replace their hypotheses with new ones, in the following way: each inactive agent **A** contacts randomly another agent; if the agent contacted is active, agent **A** copies its hypothesis, and if the agent contacted is inactive, agent **A** selects randomly a new hypothesis from the search space. The process continues until some appropriate terminating condition is met.

```
1 - All agents select hypothesis/location, randomly
2 - Loop (until some terminating condition is met)
    # Test phase – loop on all agents
      * Each agent selects and evaluates a micro-feature.
      * If OK, agent is said to be active, otherwise inactive.
    # Diffusion of information – loop on inactive agents
      * Each inactive agent selects randomly another agent.
      * If agent contacted is active, its hypothesis/locations is
        copied, otherwise a new location is randomly selected.
    End loop
```

**Table 1: pseudo-code for the standard SDS algorithm.**

During the diffusion phase, inactive agents move to different locations, i.e. select new hypotheses. In the beginning, few agents are active, and consequently most of the inactive agents explore the search space randomly. Therefore, most of the computational resources are spent on the exploration of the search space. As the process continues, more agents become active, i.e. find promising locations, which attract more agents during the diffusion phase. This is a positive feedback mechanism that ensures the eventual concentration of agents in the optimal locations (i.e. the agreement of agents on the optimal hypotheses). On the other hand, in each iteration, each active agent randomly selects a micro-feature, and performs the corresponding partial evaluation of the agent's current location; this ensures that agents are not stuck in suboptimal locations.

Standard SDS does not make any assumptions about the search space and it is essentially a random search. What makes SDS different of other generic search techniques is that agents only perform partial evaluations of the objective function. Computational resources are not wasted on unpromising candidate solutions, which makes SDS an efficient technique for search problems where the full evaluation of the objective function is computationally expensive.

## 3. SDS AS A FRAMEWORK FOR FEATURE TRACKING

SDS is essentially a generic framework to address search problems. To that framework we need to add an appropriate objective function that measures the suitability of candidate solutions to a particular search problem. The choice of such a function is domain-dependant, and when thinking of possible alternatives there are two main questions to consider:

1) In which conditions will the optimal location (which depends on the similarity function) provide a good solution to the problem in the application domain?
2) How does it affect the computational cost of the calculation of the optimal location?

The first question may look unnecessary, but it becomes clear that it is an essential point when we translate it to the application domain: in which conditions would the solution to the search problem yield a good estimate of atmospheric wind? Different objective functions, including Euclidean distance and cross correlation, may give different estimates of displacement.

In this study we have considered two objective functions to use with the SDS framework, both defined in terms of partial evaluations. In the following definitions we use **i** and **j** to denote pixels (i.e. image coordinates) and **v** to denote image displacements (i.e. all **i**, **j** and **v** are 2-dim vectors), and T with a suitable subscript to denote a partial evaluation of the similarity function. The search space is therefore 2-dimensional and a point in the search space represents a motion vector in the sequence of images.

Function F1 - a micro-feature is associated to a pixel **i** in the template box. We define the partial evaluation $T_i(\mathbf{v})$ of a displacement **v** in the following way: $T_i(\mathbf{v})$ is positive if

$$| \text{Rad1}(\mathbf{i}) - \text{Rad2}(\mathbf{i}+\mathbf{v}) | < \varepsilon,$$

for a suitable $\varepsilon > 0$, where Rad1 and Rad2 represent radiance values in the first and second images, respectively.

Function F2 - a micro-feature is associated to a pair of pixels **i**, **j** in the template box. We define the partial evaluation $T_{ij}$ of a displacement **v** in the following way: $T_{ij}(\mathbf{v})$ is positive if

$$\text{Sign} ( \text{Rad1}(\mathbf{i}) - \text{Rad1}(\mathbf{j}) ) = \text{Sign} ( \text{Rad2}(\mathbf{i}+\mathbf{v}) - \text{Rad2}(\mathbf{j}+\mathbf{v}) ).$$

A useful concept to discuss results is the test score of a location **v**. Given an objective function decomposable into micro-features, the test score of a location is the percentage of micro-features that would result in a positive evaluation of that location.

It has been mentioned earlier here that standard SDS is essentially random search and it does not make any assumptions on the nature of the similarity or correlation surface. However, in the case of AMVs these surfaces are smooth. Around the best location there are good locations, so good locations can be used to drive the search process by increasing the probability or selection of neighbour locations. Standard SDS can be modified so that the a priori knowledge of the nature of the similarity surfaces can be exploited to improve the efficiency of the search.

## 4. PRELIMINARY RESULTS

We have used WV 6.2 μm Meteosat-9 images for the tests. We chose to use this imagery in order to simplify the problem as much as possible, avoiding complications such as coastlines and multilayer situations.

For the first set of tests we produced an artificial sequence based on a real image (see figure 1). It was used as the first image in the sequence; the second image was created by applying a predefined displacement to the first (5 pixels in the East direction and 7 pixels in the South direction). This setup is not realistic, but it was very useful to test the software system developed for the experimentation. We used three types of similarity functions: the one-pixel and the two-pixel functions defined in section 3, both based on the SDS-framework, and Euclidean distance (mainly for comparison). The size of boxes was 16*16 pixels in all tests, and only boxes fulfilling the condition "contrast > 48 & standard deviation > 8" were selected as targets for matching. The selection is illustrated in the figure: the origins of vectors are the centres of those boxes fulfilling the selection conditions. The choices on box size, contrast and standard deviation are just practical values to start with the experimentation.

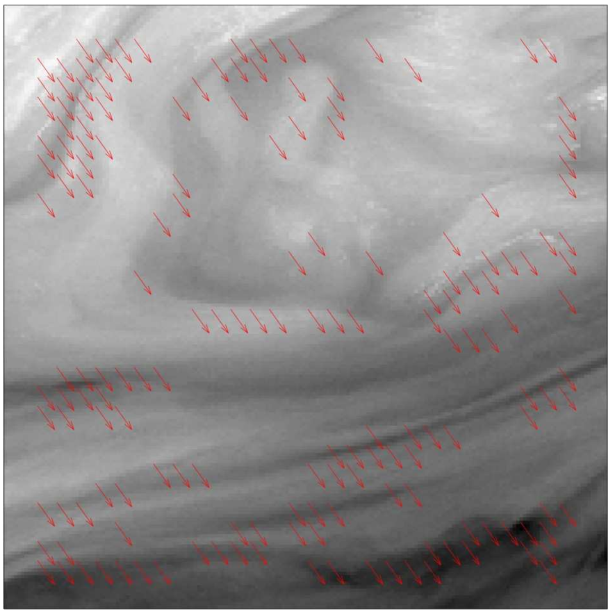

**Figure 1: Meteosat-9 WV 6.2, 17 July 2007, around 10 UTC, North Atlantic. Superimposed motion vectors calculated from an artificial unrealistic sequence.**

Although tests run with the artificial sequence were useful, the setup is not realistic. Unlike with a real sequence, there is exactly one perfect match for every target. The three objective functions produced exactly the same ("true") motion vectors for every target. Regarding computational costs, SDS-based functions were 5 to 10 times faster than Euclidean distance. However, this is not representative, partly because implementations were not optimised, but mainly because the convergence of the SDS-based methods is very fast when there is a perfect match (once an agent enters a perfect location, it stays there and attracts other agents).

A second set of tests was prepared with a real sequence of two images (the image in figure 1 and the real subsequent Meteosat-9 image, 15 minutes later). In this case, as there is no perfect match, the application of selection conditions is essential in order to obtain meaningful motion vectors. Figure 2 shows similarity surfaces, for the three functions, for a 16*16 template situated in the upper left corner of the image and satisfying the target selection conditions. The size of the search space is (-20, 20) pixels in both x and y directions. The two main points to mention, regarding the calculation of motion vectors are: 1) different similarity functions often yield different motion vectors (although they are usually very similar), and 2) convergence of SDS-based methods is slower than in the first set of tests.

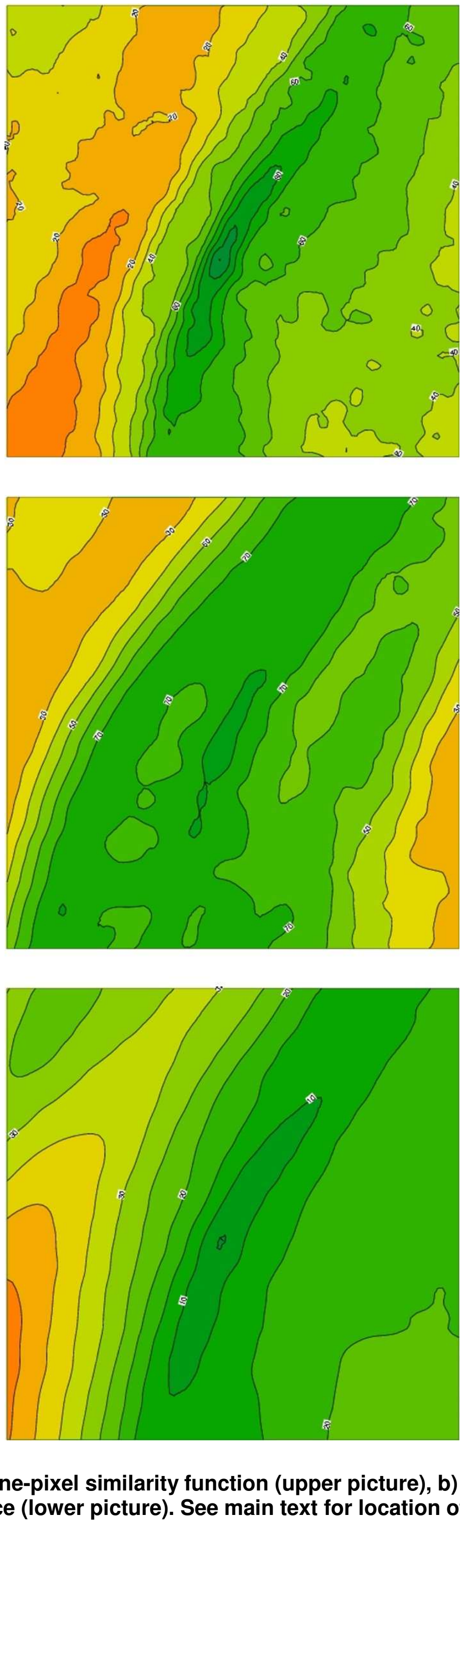

**Figure 2: similarity surfaces for a) one-pixel similarity function (upper picture), b) two-pixel similarity function (middle picture) and c) Euclidean distance (lower picture). See main text for location of target and size of search space.**

Figure 3 shows a similarity surface for a box, from the middle left part of the image (in figure 1), selected intentionally so that it does not satisfy the target selection conditions. Unlike the surfaces presented in figure 2, the area representing the best matches is flat and wide, like a plateau. Only the surface associated to the one-pixel SDS-based function is shown here, but the two other functions also showed plateau-like surfaces, although different from each other.

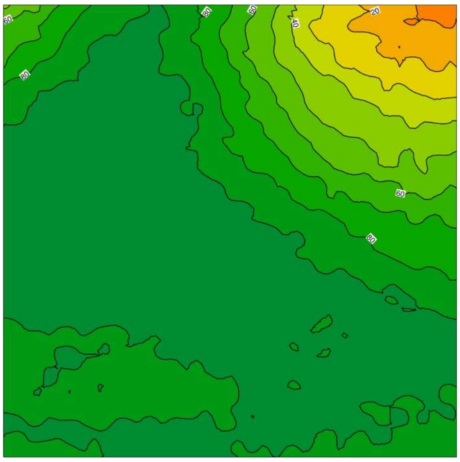

**Figure 3: similarity surfaces for the one-pixel function. See main text for location of target and size of search space.**

## 5. CONCLUDING REMARKS

This paper has presented the research carried out to explore the use of Stochastic Diffusion Search as a framework to tackle feature tracking in the context of AMV derivation. The research is in early stages, and the work done so far has brought several basic (and old) issues to the front:

- First of all, what is exactly a motion vector? Motion vectors are not uniquely determined by the sequence of images: both the box size and the function representing similarity between boxes affect the resulting motion vectors.
- In which conditions is a motion vector a good estimate of atmospheric wind? At which scale?
- How to characterise targets likely to yield, when using an appropriate similarity function, a good motion vector, i.e. good estimate of atmospheric wind?
- Given suitable targets, which objective functions are likely to yield good motion vectors? How does the choice of the function affect the computational efficiency? Can the choices of objective function and target selection procedure be independent?
- Finally, would an alternative mathematical representation of the radiance field be more appropriate? Different representations (e.g. wavelets or Fourier expansion) would allow the definition of different objective functions, perhaps more appropriate than those considered here.

Plans for the future include exploring the issues mentioned above. The research work could also be extended in other directions:

- Apart from displacements, also rotation and deformation could be included as possible transformations. An extension to 3, 4 or 5 -dimensional search spaces would be natural within the SDS framework.
- The basic search framework can be modified in such a way that the process is able to "sense" the characteristics of the similarity landscape and adjust itself to it accordingly.

## ACKNOWLEDGEMENTS

Attendance of the first author to the workshop was financially supported by Eumetsat and by the University of Reading's Research Travel Grant Sub-Committee. We thankfully acknowledge their support. We also thank Vesa Karhila for his interesting comments.

## REFERENCES

Bishop, J. M. (1989). Stochastic Searching Networks. Proceedings of the First IEE Conference on Artificial Neural Networks, London, UK.

Bishop, J. M. and P. Torr (1992). The Stochastic Search Network. In R. Linggard, D. J. Myers, and C. Nightingale (eds.), Neural Networks for Images, Speech and Natural Language. Chapman & Hall, New York.

Borde, R., and R. Oyama (2008). A direct link between feature tracking and height assignment of operational AMVs. Proceedings of the Ninth Int. Winds Workshop, Annapolis, Maryland, USA, 14-18 April 2008. Available from EUMETSAT.

Bresky, W., and J. Daniels (2006). The Feasibility of an Optical Flow Algorithm for Estimating Atmospheric Motion. Proceedings of the Eighth Int. Winds Workshop, Beijing, China, 24-28 April 2006. EUMETSAT P.47.

De Meyer, K., S. J. Nasuto, and M. Bishop (2006). Stochastic diffusion search: Partial function evaluation in swarm intelligence dynamic optimisation. In Abraham, A., Grosan, C., Ramos, V., eds.: Studies in Computational Intelligence. Volume 31. Springer Verlag, Germany (2006).

Dew, G., and K. Holmlund (2000). Investigations of Cross-Correlation and Euclidean Distance Target Matching Techniques in the MPEF Environment. Proceedings of the Fifth Int. Winds Workshop, Lorne, Australia, 28-31 March 2000. EUMETSAT EUM P.28.

Grech-Cini, J. J. and G. T. McKee (1993). Locating the mouth region in images of human faces. In P. Schecnker (ed.), Proceedings of SPIE – The International Society for Optical Engineering, Sensor Fusion VI 2059, Massachusetts.

Velden, C., J. Daniels, D. Stettner, D. Santek, J. Key, J. Dunion, K. Holmlund, G. Dengel, W. Bresky and P. Menzel (2005). Recent Innovations in Deriving Tropospheric Winds from Meteorological Satellites. Bulletin of the American Meteorological Society, 86, pp 205-223.