One dimensional Self-Organizing Maps to optimize marine patrol activities

Victor Lobo, Portuguese Naval Academy & Fernando Bação, ISEGI-UNL
ISEGI-UNL
Campus de Campolide
1070-312 Lisboa
Portugal
vlobo@isegi.unl.pt, bacao@isegi.unl.pt

Abstract – **A method for planning routes for patrol vessels is proposed. This method is based on a Self-Organizing Map (SOM) solution for the Travelling Salesman Problem (TSP), although with significant changes. The locations of reported Search and Rescue (SAR) requests, together with the locations of reported occurances of illigal fishing activities are used as guidelines for designing the path vessel should take. However, instead of forcing the patrol routes to pass exactly in those locations, as would happen in a TSP, the proposed method uses the locations as density estimators for where the patrol effort should be placed. It then obtaines a patrol route that passes through the areas with greater density. We show the behaviour of the proposed method on artificial data, and then apply the method to some data from the Portuguese Navy, obtaining possible routes for its patrol vessels.**

I. INTRODUCTION

A large number of tasks performed by the Portuguese navy are related with non-military missions, such as enforcing the protection of the Portuguese Exclusive Economic zone from illegal fishing activities, monitoring environmental risks, and search and rescue missions (SAR). The definition of the best patrol route is an interesting problem in the planning stage of these missions. Presently, this is largely based on the previous empirical experience of the officer in charge and there are no guidelines or objective criteria to support route design. Here we seek to establish a simple strategy which enables a more consistent planning of these missions, establishing an objective criteria and using a neural network (in this case, a 1-Dimensional Self-Organizing Map [1]) as the optimization tool.

We assume that the best route for a non-military mission is the route that maximizes the probability of "being useful", *i.e.*, of finding some kind of illegal activity and being close to some kind of emergency. We also assume that the locations of reported occurrences are good predictors of where incidents will occur in the future. Thus, the best route should be one that "roughly" goes through all the points where occurrences were reported recently. Using this criterion the problem becomes similar to the well known Traveling Salesman Problem (TSP) [2].

We start by pointing out the similarities and differences between our problem and the TSP. Next we give a short presentation of the SOM, it's application to the design of patrol routes, and highlight some relevant characteristics of this method. In section IV we present some experiments using artificial data, and try assess the impact of certain parametrizations. In section V we use the proposed algorithm to obtain patrol vessel routes using data from the Portuguese Navy. Finaly, we conclude with some remarks on the relevance of the proposed method, its major shortcommings, and propose future developments and improvements.

II. SIMILARITIES AND DIFFERENCES WITH THE TSP

The TSP is a well known problem that has been thoroughly studied in operations research (OR). The basic formulation can be described as follows: if a salesperson has to visit n cities, and if he desires to minimize the distance, which is the best path? Although the TSP is easily stated it has been a difficult challenge. The only way to find the optimal solution consists on calculating (at least implicitly) all the possible paths. The number of possible paths for n cities is given by $n!$. If the number of cities is small (5 or 10) it is perfectly feasible to search all possible paths. However, if the number of cities increases (for example to 500) then it becomes impossible to search all possible solutions to the problem.

The TSP belongs to the class of NP-hard problems [2], as was proved in [3]. The fundamental consequence of this is that at present the best way to solve this problem is to use heuristic methods. Although these methods do not guarantee optimal results they usually produce good solutions in adequate time. An overview of the available heuristics to solve the TSP can be found in [4].

The original TSP formulation requires that the algorithm must find a solution which passes rigorously through all the designated cities. In the case of our problem, we may relax this requirement, and only demand that the trajectory pass roughly near the occurrences points. This relaxed problem was recently addressed in [5] using algorithms which required a specific neighborhood radius to be met. We propose the use of a standard Self-Organizing Map [1, 6] as a simpler and more flexible approach to the problem.

The Self-Organizing Map (SOM) can be adapted to the TSP through the use of 1-dimensional maps. In this case the network is trained using the locations (x and y coordinates) that need to be visited. At the end of the training phase, the line representing the SOM will be laid out in such a way that it will pass in the areas where the density of data points is higher. Using the SOM to solve the TSP is not new [7-10], nevertheless all previous approaches are focused on the original TSP formulation.

Unlike in the TSP, when designing patrol routes the relative density of the points of previously reported occurrences is much more relevant than the actual position of each of the reported occurrences. Given the preemptive nature of these missions the objective is to use the historical data to loosely orientate the path, allowing for an optimal positioning of the vessel with regard to the areas where the probability of new occurrences is high. The fundamental idea is that areas of high density of reported occurrences constitute a good predictor of future occurrences.

In this paper we use the 1-dimensional SOM as an estimator of curves of maximum probability. In this procedure we substitute the original points, given by the reported occurrences, by a much smaller number of ordered SOM units. This way the SOM has to optimize the distribution of its units in order to represent a more complex reality (all the reported occurrences). In this sense this constitutes a data reduction and ordering problem where we seek to maintain the fundamental aspects of the original distribution, although drastically reducing the representational resources. The SOM units constitute the representational resources and all of them are connected to neighboring units. This means that the position of each unit influences the position of neighboring units. It is this particular feature that allows the use of the SOM in the definition of curves of maximum probability.

III. 1-DIMENSIONAL SOMS AND THEIR BEHAVIOR

Self-Organizing Maps [1, 6], or SOM for short, also known as Kohonen Neural Networks are primarily visualization and analysis tools for high dimensional data, but they have been used in many different tasks, such as clustering, dimensionality reduction, classification, sampling, vector quantization, and data-mining [1, 11]. The basic idea of a SOM is to map data patterns onto an n-dimensional grid of units (also known as neurons). That grid forms the output space, as opposed to the input space where the data patterns are. This mapping tries to preserve topological relations, *i.e.*, patterns that are close in the input space will be mapped to units that are close in the output space, and vice-versa. So as to allow an easy visualization, the output space is usually 1 or 2-dimensional.

Before training, the units are normally initialized randomly. Usually the training consists on two parts. During the first part of training, the units are "spread out", and pulled towards the general area (in the input space) where they will stay. This is called the unfolding phase of training. After this phase, the general shape of the network in the input space is defined, and we can then proceed to the fine tuning phase, where we will match the units as closely as possible to the input patterns, thus decreasing the "quantization error". The quantization error is the sum of the differences between each input pattern and its neares unit, and can be seen as a measure of how weel the units represent the input patterns. The basic SOM training algorithm can be described as follows:

Let X be the set of I training patterns $\boldsymbol{x_1}, \boldsymbol{x_2}, .. \boldsymbol{x_n}$
W be a $p \times q$ grid of units $\boldsymbol{w_{ij}}$ where i and j are their coordinates on that grid
α be the learning rate, with values in]0,1[, initialized to a given initial learning rate
r be the radius of the neighborhood function $h(\boldsymbol{w_{ij}}, \boldsymbol{w_{mn}}, r)$, initialized to a given initial radius

```
Repeat
For k=1 to n
  For all w_ij ∈ W, calculate d_ij = || x_k - w_ij ||
    Select the unit that minimizes d_ij as the winner w_winner
    Upadate each w_ij ∈ W: w_ij = w_ij + α× h(w_ij,w_mn,r)×|| x_k - w_ij ||
  Decrease the value of α and r
Until α reaches 0
```

The neighborhood function h is usually a function that decreases with the distance (in the output space) to the winning unit, and is responsible for the interactions between different units. During training, the radius of this function will usually decrease, so that each unit will become more isolated from the effects of its neighbors. The learning rate α must converge to 0 so as to guarantee convergence and stability for the SOM [1].

For our purpose the best approach is to use 1-dimensional SOMs instead of the more usual 2-dimensional SOMs. In the context of designing patrol routes the 1-dimensional SOM is particularly suited, as happens in the TSP. In both cases we attempt to map a 2-dimensional space into a 1-dimensional space (a line), i.e. a sequence of points where the patrol vessel should pass. Mappings of 1-dimensional input spaces to 1-dimensional SOMs have been extensively studied and some important properties identified by [12]. For our application one such property is of particular relevance: the SOM produces a bias in the representation of the input space. In fact, the distribution of the classification resources (units) will be more than proportional in lower density areas. This effect is usually referred to as "magnification effect" [12, 13]. The quantification of this effect in more general mappings has proved to be elusive and is still an unresolved issue. Nevertheless in our application we expect to observe the "magnification effect", which implies that even areas with a low density of reported occurances will tend to be visited.

Two issues need to be considered in the application of 1-dimensional SOMs to the optimization of patrol routes. The first regards the number of units to use, and how changes in the number of units affects the path calculation. The second issue concerns the fine-tuning of the neighborhood parameter used in the SOM training. Different neighborhood radius will yield paths with different properties. These two issues are tested in the following section.

IV. EXPERIMENTS WITH ARTIFICIAL DATA

To test the effect of different parameters on the route obtained, we generated two artificial datasets. In our experiments we produced SOMs with different numbers of units (10, 20, 50 and 100), and with different final radius for the neighborhood function (0, 1, and 2). In all experiments we used only one training phase with 100 epochs, a learning rate of 0.5 and square (or bubble) neighborhoods with an initial radius of 80% of the map size. These experiments used the SOMToolbox for Matlab implementation of SOM [14].

The first dataset consists of 500 points with a uniform distribution in a unit square, and the results are presented in Fig. 1. Two aspects of the algorithm stand out from this figure. On one hand, as the number of units increases, the route covers the input with more accuracy, winding over itself. This is an expected result as it only reflects the fact that having more representational resources the SOM will more accurately depict the density of the data points. On the other hand, as the final neighborhood increases, the path obtained will provide less detail on the distribution of the data points. To a certain extent it can be argued that when the final neighborhood reaches 0 both the unfolding and the fine tuning phase are present in the training. On the contrary, if the final neighborhood remains relatively large only the unfolding phase is present.

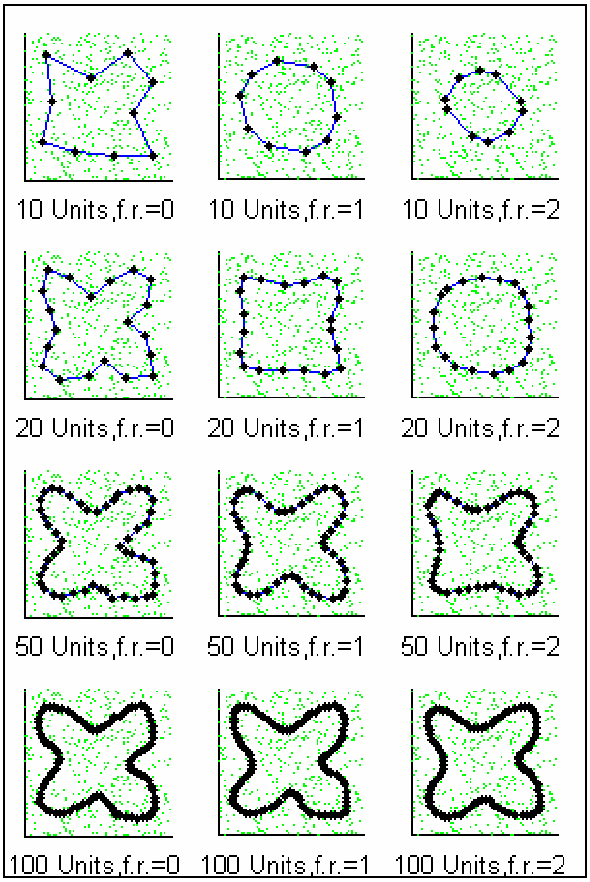

Fig. 1. 1-dimensional SOMs trained with uniform data. The black dots represent the actual location of the units.

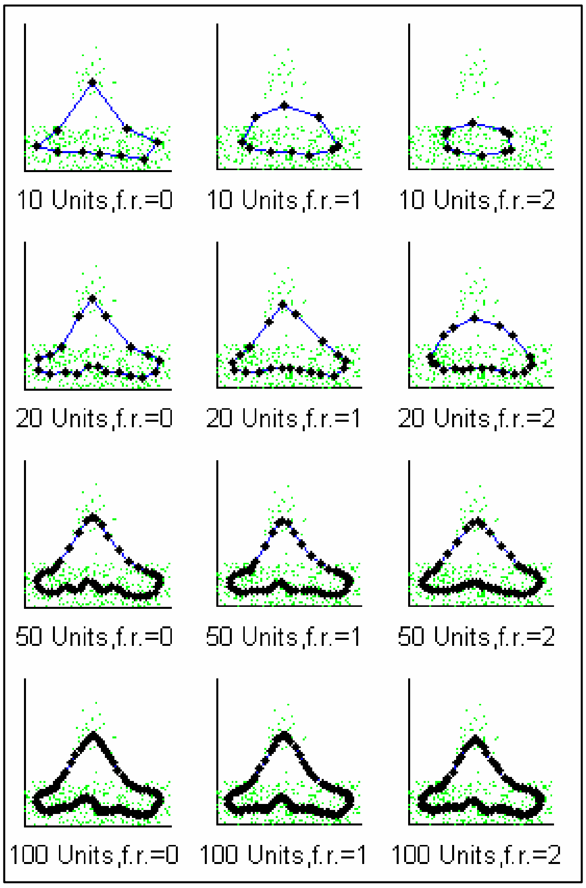

Fig. 2. 1-dimensional SOMs trained with non-uniform data. The black dots represent the actual location of the units.

The second dataset consists of 300 points with a uniform distribution in ([0 1],[0 0.3]), which roughly corresponds to the lower third of the unit square. To these points we added 30 more with a much lower but also uniform density in ([0.5 0.6],[0.6 0.7]) corresponding to a small square approximately centered in the upper half of the unit square. The results are shown in Fig. 2, and confirm the comments made above about the influence of the number of units and the size of the final neighborhood. In this particular case the presence of a small but noticeable lump of points above the main distribution allows us to confirm that, when using a similar number of units, low values for the final neighborhood will improve the representation of these outliers.

V. THE PORTUGUESE CASE

Additional tests were performed using the SOM on real world data, representing reported occurrences near the Portuguese coast line. The data used in this study was kindly provided by the Portuguese Navy. Due to confidentiality and national security issues the data provided constitutes only a sample of the occurrences. The sample comprises 93 SAR occurrences during 2003. The database file was composed of cartesian *(x,y)* geographic coordinates, derived from the actual latitude and longitude of reported occurances. The units of these coordinates are irrelevant for the purposes of this paper. We ran experiments with a large number of SOM units (twice as many as data patterns), which lead to long and “winding” routes. We also ran experiments with fewer SOM units (93, 46, and 23 units), which lead to ever shorter and “straighter” routes. In each case, we ran 50 independent tests, which due to the non-deterministic nature of SOM lead to slightly different results. Examples of the routes obtained are presented in Fig.3, and the numerical results are presented in Table 1. The average lengths obtained when using a large number of units are generally higher than the ones obtained with few units. It must be noted that although the average length obtained with 46 units is smaller than the one obtained with 23, the standard deviation is greater than the difference and thus this difference has no statistical significance. The same happens with the average lengths using 93 and 186 units. As was clear from the previous tests on artificial data, as the number of neurons decreases the path tends to disregard “outlier” occurrences concentrating on high density areas.

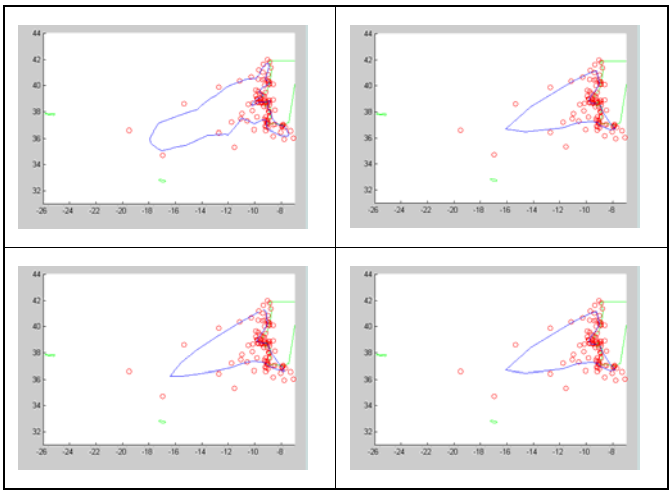

Fig. 3. Routes obtained with SAR data. In the top left route we use twice as many SOM units (186) as available data points (93), in the top right we use 93 units, in the bottom left 46, and only 23 in the bottom right.

Table 1. Average distances (and standard deviations) obtained over 50 trials, using closed paths

Nº SOM units	Average length	Standard Deviation
186	24.4130	2.5135
93	24.5072	2.6288
46	18.5052	1.2594
23	19.7034	2.3954

VI. CONCLUSIONS AND FUTURE WORK

Based on our experiments we conclude that the SOM can be used as a flexible and quick way to obtain optimized paths for patrol vessels in non-military missions. The processing time needed for the calculations is rather small (around 8 seconds) which allows real time redefinition of the path as new information is acquired. The use of locations of previously reported occurrences gives the opportunity to design a path which constitutes the informed “best guess” on where the presence of the patrol vessel is bound to discourage illegal activities and be available to rapidly attend to SAR missions. We do not, at present, show any comparisons with other methods, since no other systematic methods are used to perform this task.

Although the concept of using SOMs for this task was shown to be viable, additional work must be done to obtain good results. There are three main aspects that must be improved: forcing initial/final home ports, taking into account natural barriers, and using SOM unit density to define speed.

If we whish to specify initial and final ports for the itineraries, we may force one or two of the SOM units to have fixed coordinates, depending on whether the initial and final port are the same or not. This is very easy to do, and we did run some tests using this approach. We also conducted a series of experiments using “open end” routes, *i.e.*, routes that consitute open paths with no constains on initial and final locations. So as to make the results comparable we did not include them in this paper.

Some of the routes obtained in our tests cross land. This is obviously not possible for a patrol vessel (though it may still be useful for maritime patrol aircraft). One simple way of dealing with the problem is to sail between the two points where the route crosses the shore by the shortest sea route. This approach, while simple, does imply that we are not using the best (shortest possible) route. The correct solution to this problem would be to use true sea distances between points in the workings of the SOM algorithm. This may be achieved by calculating those distances using a Geographic Information System (GIS), which would also solve the problem of using Cartesian instead of geographic coordinates. Once again this would not affect the fundamental principals of our approach, but would require a lot of re-coding.

Finally, the density of SOM units is an indication of the density of reported occurrences. Thus, areas where the density of SOM units is greater should be patrolled more carefully. The proximity of consecutive SOM units can be used to define the speed which should be used during the patrol. However, in many cases, it is better and more economical to use a constant speed. In this case, equally spaced target-times may be given to each unit, giving the vessel’s captain the liberty to zig-zag or circle the locations of the units until those target-times are met. Additionally, there is the possibility of pruning the proposed path. Increasing distance between units indicates a decrease in point density, this way the longest edges of the SOM are primary candidates for pruning.

There is a wide range of applications for this tool. It can be used to direct unmaned areal vehicles, allowing real time path definition based on continue information feeding. It can also be used in police and military patrolling activities, where the objective is to control areas where the probability of certain events is higher than a certain threshold.

.

REFERENCES

[1] T. Kohonen, Self-Organizing Maps, 3rd ed. Berlin-Heidelberg: Springer, 2001.

[2] S. A. Cook, "An overview of computational compexity," Communications of the ACM, vol. 26, pp. 400-408, 1983.

[3] C. H. Papadimitriou and K. Steinglitz, "Some Complexity Results for the Traveling Salesman Problems," presented at Annual ACM Symposium on Theory of Computing - Proceedings of the eighth annual ACM, Hershey, Pennsylvania, United States, 1976.

[4] S. Arora, "Polynomial Time Approximation Schemes for Euclidean Traveling Salesman and Other Geometric Problems," Journal of the ACM, vol. 45, pp. 753–782, 1998.

[5] A. Dumitrescu and J. S. B. Mitchell, "Approximation algorithms for TSP with neighborhoods in the plane," Journal of Algorithms (Elsevier), vol. 48, pp. 135-159, 2003.

[6] T. Kohonen, "Clustering, Taxonomy, and Topological Maps of Patterns," presented at Proceedings of the 6th International Conference on Pattern Recognition, 1982.

[7] G. J. Hueter, "Solution of the traveling salesman problem with an adaptive ring," presented at Neural Networks, 1988., IEEE International Conference on,, 1988.

[8] C. S.-T. Choy and W.-C. Siu, "New Approach for Solving the Travelling Salesman Problem using Self-Organizing Learning," presented at IEEE International Conference on Neural Networks, 1995.

[9] H. Sasamura, R. Ohta, and T. Saito, "A simple learning agorithm for growing ring SOM and it's application to TSP," presented at 9th International Conference on Neural Information Processing (ICONIP'02), 2002.

[10] S.-J. Kim, J.-H. Kim, and H.-M. Choi, "An efficient Algorithm for Traveling Salesman Problems based on Self-Organizing Feature Maps," presented at IEEE International Conference on Fuzzy Systems, 1993.

[11] J. Vesanto and E. Alhoniemi, "Clustering of the Self-Organizing Map," IEEE Transactions on Neural Networks, vol. 11, pp. 586-600, 2000.

[12] M. Cottrell, J. C. Fort, and G. Pages, "Theoretrical Aspects of the SOM algorithm," Neurocomputing, Elsevier, vol. 21, pp. 119-138, 1998.

[13] J. C. Claussen, "Winner-relaxing and winner-enhancing Kohonen maps: Maximal mutual information from enhancing the winner," Complexity (Wiley), vol. 8, pp. 15 - 22, 2003.

[14] J. Vesanto, J. Himberg, E. Alhoniemi, and J. Parhankangas, "SOM Toolbox for Matlab 5," Helsinki University of Technology, Espoo A57, April 2000 2000.