

Evolutionary simulation of an agent based mobility system using indirect communication

Francesco Zanlungo
Department of Physics
University of Bologna
Bologna, via Irnerio 46 40126

Takaya Arita
Graduate School of Information Science
Nagoya University
Chikusa-ku, Nagoya, Aichi, 464-8601 Japan

Abstract

We study the evolutionary dynamics of a mobility system in which agents are able to explore the environment and to communicate between them in order to increase the efficiency of the system, i.e. to minimize the average time needed to reach their goal, choosing the quickest path and avoiding the formation of traffic jams. The agents use an evolvable pheromone-like indirect communication system, of which we also give an interpretation as the average of different kinds of direct communication.

Key words: Indirect communication, evolution of neural networks, mobility systems, multi-agent system

1 Introduction

A town is a complex system in constant evolution. Both the environment (the structure of the town, its dimensions, the road and public transportation system, the location of the places of interest) and the behaviour of the inhabitants are constantly changing, and influencing each other. Focusing only on the mobility system, new roads, railways and subway lines are built according to the needs of the citizens, but the presence of these new structures modifies the behaviour and habits of the citizens, and so on.

In this paper we propose a simple system of agents whose behaviour is regulated by evolving neural networks. These agents move on a 2D manhattan grid that represents a town's road system, with the connected traffic problems. The knowledge of the structure of this road system is obtained through the exploration of the system by the agents, and communicated to other agents through a "pheromone-like" field.

It is our opinion that the model that we propose can be a fruitful environment to study the evolutionary dynamics of a system of different interacting agents regulated by a simple behavioural mechanism, and eventually the emergence of some kind of "global behaviour". In particular, we are interested in studying the emergence and evolution of an indirect communication system. This model could also be useful to get some insight about the evolution of actual mobility and urban

systems.

It can also be seen as an attempt to understand if some kind of pheromone-like communication could be used to improve the efficiency of a real mobility system and avoid the formation of traffic jams.

2 The model

2.1 The mobility system

The agents are located on the nodes of a manhattan 2D network ("a town"). At each time step they have a probability $p \leq 1$ to perform just one of five possible actions, i.e. to move to one of the four neighbour sites (at manhattan distance $d = 1$, where the distance between the points P_1 and P_2 is defined by $d(P_1, P_2) \equiv |x_1 - x_2| + |y_1 - y_2|$) or to rest on the site they are located in. We want our town to have wider roads, in which higher speed is possible, and narrower roads. We also want to have a traffic problem, i.e. agents should be slowed down when a great number of them is located in the same site (when we talk about the agent's speed, we are actually talking about its probability p to perform an action). To achieve this we define, identifying a site with its discrete coordinates (i, j) and the number of agents on that site with $n(i, j)$, the probability to act as

$$p((i, j), n(i, j)) = w(i, j)t_{(i, j)}(n(i, j)) \quad (1)$$

where $0 \leq w(i, j) \leq 1$ is the wideness of the site and $t_{(i, j)}$ is a (eventually) site dependent "traffic jam function", which we ask to be (not necessarily strictly) decreasing and to have $t_{(i, j)}(1) = 1 \forall (i, j)$. In particular, we are going to use $t(n) = (1/n)^\gamma$, with $\gamma \geq 0$.

Even if we talk about "wide and narrow road" it is clear from the definition that the wideness is related to the site, and not to the roads, since no road is actually present in this model. We order the wide sites in rows inside the grid (see figure 1 for the basic structure of the town) to simulate the presence of a road, but p has no dependence on the direction of the motion so there is no difference between an agent that is following a "road" and an agent that is just crossing it.




Figure 1: Basic structure of the town’s road system. Thick means wide.

This is a very simple and abstract model that is, according to our opinion, feasible to study the general features of the interaction of a group of agents in a mobility system. Nevertheless in building the model we were inspired by the actual car traffic systems. A node in the network represents a crossroad (plus some portions of the roads leading to it). If a large number of cars is located in the same crossroad, a traffic jam occurs. w stands for the maximum speed that the car can have on that (cross)road, while t is the “tendency” of the crossroad to the formation of traffic jams.

We can also think of this model as a system of pedestrians, in which the wide roads correspond to public transportation, and t is the waiting time to use it.

2.2 Communication

The way the citizens of a town learn to use its mobility system is obviously quite complicated and differentiated (past experience, communication with other agents, centralized communication like news from the radio, etc.). In our model we wanted to introduce some kind of mechanism that could represent an average of all these processes, that could be regulated by the agents themselves and that could be utilized just through a local observation of the environment (to avoid a too complex sensory system or the introduction of memory).

We decided to use pheromone based communication which has shown to be an efficient way to evolve communities of collaborating agents [1].

We introduce a pheromone field $P(i, j)$, generated by the same agents which drop an amount q of pheromone every time that they cannot perform an action, i.e. every time that “they have to wait” because they are located in a narrow road or they are stuck in a traffic jam. This means that if the indirect communication system emerges, pheromone should be interpreted as “repulsive” and used to inform the other agents about “bad places”. Pheromone evaporates according to a coefficient c_e and turns into evaporated pheromone

$P_e(i, j)$ that diffuses with coefficient c_d [1]. Agents are able to perceive the gradient of this field, i.e. the difference with the neighbouring sites $P_e(i, j) - P_e(neigh)$.

We could interpret this pheromone field as a “rumor” that stands for the “collective knowledge” of the urban structure, and spreads between agents. Actually in our model the rumor is spreading on the urban structure, but we can think that in this model are present immobile and mobile agents, and the rumor is spreading between the (not explicitly represented) first ones. Adopting this point of view, the medium that mobile agents are using to communicate are the immobile agents.

2.3 The decision mechanism and its evolution

The evolution of neural networks through genetic algorithms has shown to be an efficient way to develop agents and robots able to move in complex environments [2].

We use a fully connected neural network with 6 inputs (the distance from the goal in x and y coordinates ($g_x - x$) and ($g_y - y$) plus the 4 “pheromone gradients”), a layer with h hidden neurons and 5 outputs corresponding to the 5 possible actions. The action with the highest output value is performed. The connections of the networks are kept fixed during the whole generation, i.e. there is no learning, just evolution of connection weights and of the number of the hidden neurons (h).

We define the fitness function in the following way. Each agent is created in a randomly selected starting point (s_x, s_y) . If τ is the time to reach its randomly selected goal point (g_x, g_y) (a minimum distance between start and goal is imposed), the ratio $r = (|g_x - s_x| + |g_y - s_y|)/\tau$ is measured. If the agent does not reach the goal in a maximum time T , r is set to zero. The operation is repeated R times and the average value of the ratio is the agent’s fitness $f = \bar{r}$.

Since we did not know the value of h that could solve our problem, we used the following evolution strategy, (inspired by [3, 4]) to optimize also this parameter. We start using different networks divided in N_s species, according to their structure (i.e. to the value of h), each one with an equal number of members (in each experiment we have networks with different structure and weights interacting and competing at the same time on the urban structure). Tournament selection is performed inside a single species for g_s generations, then the dimension (number of agents) of the species is changed according to its fitness. The size of the tournament is changed according to the species’s new size, and evolution is performed again for g_s generations inside the species’s boundary, and so on. Mutation of the connection weights was performed adding, with probability $p_m = 0.05$, a gaussian noise with mean zero .




Figure 2: Evolution of the average fitness value through generations in the case of no traffic jams.

3 Experiments

We used 4000 agents and a 50x50 grid, averaging over $R = 20$ runs. Wide roads were defined to have $w = W = 1$ and narrow ones to have $w = N = 0.05$. The values of the pheromone parameters were $q = 0.1$, $c_e = 0.1$, $c_d = 0.1$ (we reserve for future works the task of optimizing these parameters through evolution, and eventually also of evolving the pheromone dropping mechanism).

3.1 No traffic jams

In a first class of experiments we used $\gamma = 0$ (i.e. $t(n) = 1 \forall n$) on the whole grid. In this case no traffic jam can be generated, and the only problem is to find the quickest path to the goal. We first fixed $w(i, j) = N \forall (i, j)$ and checked that we were able to evolve networks with the maximum possible fitness $f = 0.05$. Then, without introducing pheromone, we used a structure of wide roads according to figure 1. The networks evolved to $f = 0.064$ which we assumed to be the maximum fitness in absence of further information (the fitness attained by agents going straight to the goal). We will consider pheromone communication to be successful if it is able to improve this value.

The agents that dropped pheromone evolved to reach a fitness roughly 3 times higher ($f = 0.202$), i.e., indirect communication had emerged. In this experiment, pheromone communication is used to mark the narrow roads (wide roads correspond to the minima of the pheromone field). The agents developed an ability to reach the wide roads following the decreasing pheromone gradient when they were distant from their goal, to move on the the minima (follow the wide road) in order to approximate the goal, and then to leave the minima (despite the positive pheromone gradient) in order to reach the goal when they were near to it. The results of this experiment are shown in figures 2 and 3.




Figure 3: Evolution of the h species' size for the same simulation shown in figure 2. Initially 10 species (corresponding to values of h between 5 and 14) are present, then $h = 9$ (green) and $h = 14$ (red) invade the population.

3.2 Traffic jams in a uniform grid

We then fixed $w = W$ on the whole grid and used $\gamma > 0$, i.e. the agents did not have to find any “best path” but just to avoid traffic jams.

Using $\gamma = 1$ agents going straight to the goal performed with an average fitness $f = 0.292$, while the pheromone dropping agents had an higher fitness ($f = 0.377$). We tested the performance of these agents without allowing them to drop pheromone and we unexpectedly found that their fitness did not change at all. Actually the genetic algorithm had developed a network that reached the goal moving always clockwise (in the specific case). Due to the nature of the manhattan metric this allowed them to reach the goal on a minimal distance path, reducing strongly the probability of a collision (through the introduction of a “circulation rule”, i.e. of some kind of collective behaviour, as clock-wise or counter clock-wise circulation, see figure 4). The emergence of this rule has no relationship with communication, since it was based only on the “distance to the goal” inputs. Actually, for this low value of γ , the circulation rule was enough effective to avoid the formation of traffic jams, and it prevented the emergence of communication.

In the case of $\gamma = 2$ the agents going straight to the goal without using a circulation rule performed with $f = 0.024$, while the pheromone dropping agents reached a (4 times higher) value of $f = 0.098$. If we made them run without pheromone, i.e. without communication, they performed with $f = 0.030$. This value is quite lower than the one they had when we allowed them to communicate, but higher than the one attained by agents going straight to the goal. This shows that both communication and a circulation rule had emerged.

We tested on this experiment also the pheromone dropping agents that we had evolved in the previous experiment without traffic jams. Despite the completely




Figure 4: An agent has to reach point B from A, and another A from B. Black corresponds to a circulation rule that avoids collision, while red corresponds to go straightly to the goal. Notice that the two paths cover the same distance in manhattan metric.

different nature of this experiment, they were able to double the fitness of the agents going straightly to the goal ($f = 0.051$). This happens because the agents, trained in the first experiment to follow the wide roads as minima of the pheromone field, result also able to avoid the traffic jams (maxima of the pheromone field) in this experiment.

3.3 Complete model

In real mobility systems, the greatest traffic problems occur when one of the main communication routes is blocked due to a traffic jam (often caused by a traffic accident). To simulate this problem on our simple geometry we used the basic structure of figure 1 with a traffic function that used $\gamma = 1.75$ for the sites with $w = W$ and $\gamma = 0$ where $w = N$ (this can be interpreted as a rule that says that accidents, and thus traffic jams, are easier to occur where the speed is high). In this way we created a dilemma about the use of wide roads.

In this environment, the agents going straightly to the goal performed with values that changed strongly according to different repetitions of the experiment, but that were always smaller than 0.05 (between 0.035 and 0.048 in the experiments that we have performed). The crucial point is that these values are lower than 0.05, i.e. the introduction of wide roads on which traffic jams are easy to occur is harmful if they are not used in a “intelligent” way.

We then tested with these conditions the network evolved dropping pheromone in absence of traffic jams, which performed with $f = 0.063$ (almost the same value on each test), i.e. it was able to use the wide roads in a way good enough to improve over the value of 0.05. Furthermore, if we added to the road system a further “ring” of wide roads, the fitness of the agents going straightly to the goal decreased again (0.033-0.037),

while the one of agents using pheromone increased to 0.067.

We finally tried to evolve explicitly a network for this problem. The maximum reached average fitness was 0.092, but this fitness dropped to 0.060 when we tried to extract a “best” uniform population.

4 Conclusions

We have shown that simple neural networks are able to evolve a (pheromone based) indirect communication system that is effective in solving both the problem of finding a “good path” and to avoid “traffic jams” in a simple mobility system.

We have found that a network evolved in a time-constant environment (no traffic jams) could be effective also in the other two experiments without further evolution, showing the effectiveness of the pheromone communication in both the tasks of avoiding problems (pheromone maxima) and following good paths (minima). The cases with traffic jams (i.e. the environment is constantly changing according to the behaviour of the agents) showed the emergence of global behaviours (“circulation rules”) and the presence of a not trivial evolutionary dynamics (the presence of differentiated neural networks, different in connection weights and structure, seems to cause an higher value of the fitness), which should be further investigated in future works.

References

- [1] Y. Nakamichi and T. Arita, “Effectiveness of Emerged Pheromone Communication in an Ant Foraging Model”, Proceedings of the Tenth International Symposium on Artificial Life and Robotics, 324-327, 2005.
- [2] S. Nolfi and D. Floreano, “Evolutionary Robotics”, MIT Press, 2000.
- [3] K. O. Stanley and R. Miikkulainen, “Evolving Neural Networks through Augmenting Topologies”, Evolutionary Computation 10(2), 99-127, 2002.
- [4] F. Gomez and R. Miikkulainen, “Solving Non-Markovian Control Tasks with Neuroevolution”, Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, pages 1356-1361, Morgan Kaufmann, San Francisco, California, 1999.