

Simulated Road Following Using Neuroevolution

Aparajit Narayan^(✉), Elio Tuci, and Frédéric Labrosse

Aberystwyth University, LLandinam Building, Aberystwyth SY23 3DB, UK
`{apn3,elt7,ffl}@aber.ac.uk`

Abstract. This paper describes a methodology wherein genetic algorithms were used to evolve neural network controllers for application in automatic road driving. The simulated controllers were capable of dynamically varying the mixture of colour components in the input image to ensure the ability to perform well across the entire range of possible environments. During the evolution phase, they were evaluated in a set of environments carefully designed to encourage the development of flexible and general-purpose solutions. Successfully evolved controllers were capable of navigating simulated roads across challenging test environments, each with different geometric and colour distribution properties. These controllers proved to be more robust and adaptable compared to the previous work done using this evolutionary approach. This was due to their improved dynamic colour perception capabilities, as they were now able to demonstrate feature extraction in three (red, green and blue) colour channels.

Keywords: Road-following · Genetic algorithm · Neural network · Dynamic dimensionality reduction · Autonomous navigation · Active vision

1 Introduction

Autonomous navigation in its entirety is a vast and diverse field of study and research tends to be focussed on a number of sub-areas, such as steering control, obstacle avoidance, road-following, power management and road-sign detection. Amongst these, road-following or automatic driving on roads is an essential foundation of any system with desired autonomous navigation capabilities. While it may seem a trivial problem from a human perspective, accurately extracting the desired features in the environment and using them to navigate the road successfully is indeed a significant problem in terms of an AI system. This is particularly due to the amount of variance and non-uniformity present in terms of the geometry and colour composition of the road/non-road surfaces. Weather conditions such as rain, shadows, changing sunlight, etc., all have an effect on the systems visual perception of the environment and further complicate this problem.

1.1 Related Work

The design philosophy behind most engineered road-following solutions is based on maintaining an internal model of the road/non-road environment which is continuously updated based on the features extracted from the world [7]. A commonly used technique is to use sensor fusion, combining sensory data from multiple cameras and laser range-finders, to produce a more detailed and accurate representation of the world. This approach to the road following problem, especially when adaptive techniques for maintaining the model of the environment are used has been successful in real-world trials. However at the core of most such hand-crafted controllers is the issue of designer bias and the assumptions that are made of the road with regards to its geometry, contrast and colour composition. Thus successful performance may be guaranteed in environments accounted for in the design process, but often not across the entire range of possible scenarios such as in the case of [3] and [6] where geometric assumptions and limited detail meant that the model was less suitable for more complex road-shapes. There have indeed been a number of AI vehicles capable of complete autonomous navigation over the years. The foremost of these in recent times is Google's driver-less car project which logged over 500,000 km accident free during its road-testing phase. Others include Stanford's Stanley AI vehicle [8], which won the 2005 DARPA grand-challenge after successfully completing a challenging unstructured off-road course of 212 km. However in the case of such systems the cost and hardware requirements often make their implementation prohibitive in smaller low-power platforms. There have been a few attempts to use traditional machine learning strategies to train neural networks to provide full navigation control or at-least lateral steering control for autonomous vehicles, the foremost among them being the ALVINN project [2]. The neural network employed was a three layer feed-forward architecture with a single feedback unit. The input layer was fed in readings from camera pixels and a laser range-finder. This initial road-following controller paved the way for the ALVINN-VC [7] which was a more complete road-navigation system capable of dealing with junctions and intersections.

One of the key challenges of the project was to provide data for the back-propagation algorithm to train the network. In the case of road-following, training on the basis of real-world conditions to account for all the variations in the road/non-road environment would be logistically impossible. Therefore great effort was taken to create a simulated road-generator which would supply images based on the variations of as many as 200 parameters. Later trials involved training the network on sensor and motor inputs generated by an actual human driver in control. The main issue with the back-propagation approach to learning in general is over-fitting to the training data and thus rendering the system less effective in new un-encountered environments. Moreover there was still a level of human bias manifesting in the choice and generation of the training environment as well as the dictating of what the desired or perfect driving output of the controller should be. Such a control system, trained on human-driving data would never be able outperform a human driving system and its best case scenario is

that of matching the human driving. It would also not account for unexpected scenarios such as the ability to recover from steering errors and deviations.

The evolutionary machine learning approach outlined in this paper attempts to provide an alternate solution to the road-following problem, one with minimal hardware/computational requirements yet with enough adaptivity to solve the entire range of road-following scenarios. It attempts at further reducing the dependency on human-foresight and allowing the AI control system to be in charge of learning its own feature extraction and control strategies. The authors of [1] who first implemented this approach, made use of a neural network with architecture similar to the one used in [2], but instead of supervised learning the authors used evolutionary computation. Apart from having outputs for controlling motor actions, the network had a further three outputs which were fed back to the input layer and were capable of influencing the perception of the input image. Instead of having separate modules for action and perception, the paper proposed a unified motor-sensory unit. This model bears similarity to the learning methodologies of biological organisms where functional behaviour is developed through interactions with the environment and a clear link is present between actions and their effect on the perception of the scene.

Thus the aim was to evolve a controller capable of road following behaviour with the ability to dynamically change its perception of the road as needed. Because of the prohibitive logistics of carrying out the learning on real-world platform, the evolution needed to take place in a simulated environment with the option of later transferring a successfully evolved controller to a real-world platform. As an initial proof of concept the experiment was successful in showing that such controllers can indeed be evolved to successfully carry out road following across a number of simulated environments. However there were limitations with regards to their dynamic colour perception abilities and as a result their performance in certain types of scenes which they had not experienced during evolution. This paper details further progress of solving the road-following problem in simulated environments using this active vision evolutionary robotics approach and aims at addressing the limitations of the previous methodology, techniques to ensure increased robustness and adaptability of the evolved neural networks, as well as further analysing and evaluating their behaviour. It is hypothesized that the strategies outlined in this paper would enable the evolution of controllers which would be capable of ultimately performing in real-world poorly delineated and unstructured roads.

2 Neural Network Controller

A Continuous Time Recurrent Neural Network (CTRNN) is used to control the robot as shown in Fig. 1. Eqs. 1, 2, and 3 define the activation values for the 25 input, 6 hidden and 7 output neurons. In these equations, y_i represents the cell-potential, τ_i the decay constant, g the gain factor, I_i the activation of the i^{th} sensor neuron, w_{ji} the weight of synaptic connection from neuron j to neuron i , β_j the bias term and $\sigma(y_j + \beta_j)$ the firing rate. All input neurons share the same

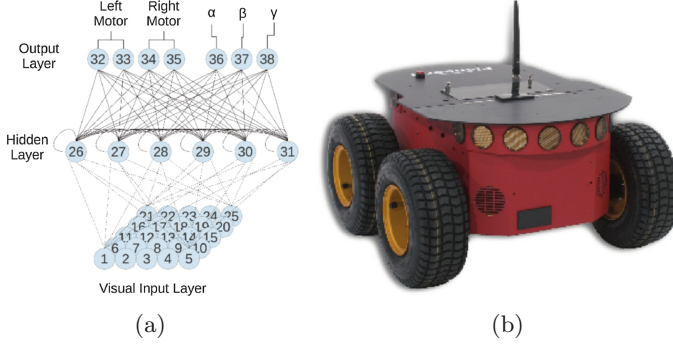


Fig. 1. (a) Architecture of the neural network controller. (b) The Pioneer robot.

bias (β^I); the same being true for output neurons (β^O). $\sigma(x) = (1 + e^{-x})^{-1}$ is the sigmoid function. The decay constants, bias terms, weights and gain factor are all genetically specified network parameters.

$$y_i = gI_i; i \in \{1, \dots, 25\} \quad (1)$$

$$\tau_i \dot{y}_i = -y_i + \sum_{j=1}^{j=31} \omega_{ji} \sigma(y_j + \beta_j); i \in \{26, \dots, 31\} \quad (2)$$

$$y_i = \sum_{j=26}^{j=31} \omega_{ji} \sigma(y_j + \beta_j); i \in \{32, \dots, 38\} \quad (3)$$

Due to the computational overheads associated with updating neural networks with large input layers, the number of input neurons was limited to 25. The image from the camera is divided into 25 equal-sized blocks. For each block, we compute the averaged red (\bar{R}), green (\bar{G}) and blue (\bar{B}) (i.e., average pixel value). Each block is associated with an input neuron and the final value I_i fed into an input neuron is computed in the following: $I_i = \alpha \bar{R} + \beta \bar{G} + \gamma \bar{B}$. The parameters α , β , and γ are generated by the network at each updating cycle, and normalised such that $\alpha + \beta + \gamma = 1$. These parameters give the system its dynamic dimensionality reduction properties. Each output neuron can increase or decrease the magnitude of these parameters to enhance or diminish the colour channel it is associated with, while at the same time having the opposite effect on the other two channels. For example, in an environment where red is the channel which shows contrast between road and non-road, having α at a maximum and the other channels at a minimum would enable the network to be presented with the best possible contrast from the scene. Figure 2 shows this effect of enhancing the correct colour channel to produce contrast between inputs corresponding to road and non-road areas.

The motion control is based on the 2D two-wheeled differential drive kinematics model for mobile robots detailed in [9]. This model takes into account the robots structural parameters i.e. radius, wheel distance and speed-limits to

give an output in terms of the robots updated position and orientation. The output of neuron 32 to 35 (Fig. 1) are used to set the left and the right wheel speeds. Complex dynamical properties such as friction are not accounted for in this model. The author in [4] highlights examples of the successful portability of this model from simulated to real-world platforms.

3 Genetic Algorithm

A population size of 52 individual chromosomes is used, with a generational limit of 3000. Trials involved the network controllers trying to perform road-following in either six or twelve simulated environments. The best individual of each generation is guaranteed a place in the next generation, whereas the one which performed the worst is truncated and made unavailable for breeding. The rest 51 individuals of the new generation are generated by breeding with the parent chromosomes selected using the roulette-wheel method. Crossover and mutation probabilities are set at 50 % and 5 % respectively. These operators remain static and non-adaptive throughout the evolution. Carrying out this process of artificial evolution over 3000 generations in a sequential process would mean an unreasonably high training time. Thus the genetic algorithm is parallelized using MPI and implemented on the HPC Wales computing cluster. Each individual runs its evaluations as a separate process and the respective fitness values are communicated to a root process which in turn carries out the evolution, generating the new generation of controllers.

4 Simulation Scenes

The evaluation scenes are the virtual environment where each controller (i.e., chromosome) is evaluated. These scenes form the basis for the network’s learning process, and the importance of this aspect needs to be stressed. These scenes have been designed to facilitate the evolution of dynamic colour perception strategies (i.e., the adaptive variation of α , β , and γ). The evolution scene graphics (see Fig. 2a) are rendered using OpenGL and are designed to simulate a camera pointing down at the ground such that the road and surroundings on either side are visible till a vanishing point further away.

The road is rendered using a modified version of the road generation algorithm employed in [1]. A total of 11 tiles are used each 160 cm long and 100 cm wide. The length of the road the robot needs to travel is 17.6 m. The virtual robot model has a diameter of approximately 54 cm. The road starts off with a smooth bend; each tile rotated 30° left or right. The direction of this turn alternates for consecutive trials. This is followed by a similar smooth bend, with greater probability (6/7) of it being in the opposite direction as the first one. This provision allows a controller to demonstrate the ability to make both kinds of turns and ensures the robot needs to be constantly maintaining its course to stay on the road. Subsequent turns are random, but checks are made to ensure no unrealistic or intersected road shapes are generated. The scene in each trial varies in

Table 1. Colour combinations of the twelve evaluation scenes.

Scene	Road	Non road	Random (Noise)
1	Bright Blue	Dark Blue	Red and Green
2	Bright Green	Dark Green	Blue and Red
3	Bright Red	Dark Red	Blue and Green
4	Bright Red, Dark Green	Dark Red, Bright Green	Blue
5	Bright Blue, Dark Red	Dark Blue, Bright Red	Green
6	Bright Green, Dark Blue	Dark Green, Bright Blue	Red
7	Dark Blue	Bright Blue	Red and Green
8	Dark Green	Bright Green	Blue and Red
9	Dark Red	Bright Red	Blue and Green
10	Dark Red, Bright Green	Bright Red, Dark Green	Blue
11	Dark Blue, Bright Red	Bright Blue, Dark Red	Green
12	Dark Green, Bright Blue	Bright Green, Dark Blue	Red

terms of the colour of the road and non-road surfaces as shown in Table 1. These scenes are created such that no contrast can be perceived between the road and non-road surfaces unless the robot is able to vary the value of α , β , and γ in an adaptive way. The 12 scenes can appear in three different formats, which differ in terms of the intensity difference between the dark and the bright colours (see Table 2).

To simulate the effect of poorly delineated roads, the edges of the textures were blended together such that there would not be a clear demarcating line between the road and non-road areas. It should be noted however that evolutionary runs carried out in roads without this effect (i.e. having a clear edge) did not demonstrate any behavioural difference. This can be attributed to the extremely low resolution of the final input image (25 pixels), which causes the network to be immune to such minor environmental variations. An additional road tile with higher levels of delineation and uneven geometry was created to be used in the testing period to assess the robustness of the evolved controllers.

Table 2. Contrast and colour distribution characteristics for the three sets of scenes.

Set	Contrast between mean intensities of road and non-road (0–255)	Range of distribution of intensities (0–255)
A	120 for all scenes	120 for all scenes
B	150 for mono-colour, 120 for dual-colour	10 for mono-colour, 30 for dual-colour
C	80 for all scenes	80 for all scenes

5 Road Bounds Checking and Fitness Function

Each trial is allowed a maximum of 250 iterations with a check being carried out after the end of each iteration (update) to see if the robot is still on the road. If the robot is detected to have moved off the road, the trial gets terminated. At the end of each trial the distance travelled is calculated by the number of road-tiles traversed thus far and the position in the current tile. In case of the trial being terminated due to the robot going off the road the current score value is divided by 5, to make the contribution of progress in the current tile negligible. This distance value $d(e)$ for each evaluation is further normalized to the range of 0.5–1.0 to present the final product, which would otherwise be a result of the powers of twelve or six, in an acceptable range. The final fitness function (Eq. 4) comprises of two components multiplied with each other, the product of distance values of each evaluation and the other a colour term Δ . In initial experiments, it was observed that the best individuals in the early stages of evolution were able to solve only a subset of the 12 scenes. These individuals dominated the population over generations, resulting in local maxima wherein the ability to solve the other scenes did not evolve. This happened in the case when the fitness was determined simply by the average distance value across all the trials. Thus having the fitness comprising of the individual distance values multiplied with each other ensures that such skewed solutions cannot dominate the population disproportionately and only individuals which perform consistently well in all the scenes are rewarded. Furthermore the Δ term was introduced to aid or guide the final solution by rewarding the correct activation of the colour outputs in each of the evaluation scenes. Populations initialized with the same random seed were tested in evolutionary runs with and without this colour term Δ to study its effect, and successful evolution was observed only in those runs where it was included.

$$F = \Delta_{final} \times \frac{1}{E} \prod_{e=1}^E (0.5 + (\frac{d(e)}{22})); \quad (4)$$

$$d(e) = NT + CS \quad (5)$$

$$CS = TL - \mu; \quad (6)$$

$$\Delta_{final} = \frac{1}{E} \sum_{t=1}^{t=E} C(e); \quad (7)$$

$$C_{1,2,3,7,8,9} = \sum_{s=50}^{s=S} |OR_s - OW_s^1| + |OR_s - OW_s^2| \quad (8)$$

$$C_{4,5,6,10,11,12} = \sum_{s=50}^{s=S} 2 \times OW_s \quad (9)$$

with $E = 12$ being the total number of trials; NT equal to the number of tiles crossed; CS equal to the score on the current tile; TL equal to the tile length; μ equal to the length of the error vector from the mid-point of the end of the




Fig. 2. (a) Grayscale image of scene 10 (see Table 1) (b) Pixel averaged version of the image with the three colour outputs (α , β and γ) equal. No contrast visible between road and non-road pixel grids. (c) Final pixel averaged version of the image with β activated, showing contrast between grids corresponding to the road and non-road areas (Color figure online).

road tile to the current position of the robot; $C(e)$ corresponding to the quality of the dynamic color perception strategy in trial e ; OR_s being the value of the colour parameters (i.e., α , β , or γ) that has to be used to discriminate between road and non-road; OW_s^1 and OW_s^2 being the values of the colour parameters (i.e., a combination of α , β , and γ) that do not discriminate between road and non-road in mono-colour scenes; OW_s being the value of the colour parameter (i.e., α , β , or γ) that does not discriminate between road and non-road in the dual-colour scenes.

A final effect of Δ_{final} is that since it gets calculated only after the 50th iteration to allow the controller time to settle on a steady sequence of colour output values for the trial, any individuals leaving the road before the 50th iteration will get a 0 for the colour score of that trial. Thus those individuals which leave the road before the 50th iteration for all the trials receive 0 as the final fitness value irrespective of any distance values gained.

6 Results and Observations

The first round of evolutionary runs was done with six scenes. These constituted of three mono-colour (1, 2, 3) and three dual-colour (4, 5, 6) scenes. Scenes were created with textures chosen from Set A (see Sect. 4). Based on the results of this stage, the experiment was extended to all 12 scenes using textures from Sets A, B and C. Each experimental condition was tested with a set of 10 random seeds, resulting in a total of 40 evolutionary runs. Due to the nature of genetic algorithms and the complexity of the problem, not all experimental runs were able to evolve a successful solution. Only those experimental runs with fitness values high enough to indicate the ability to solve more than half of the evaluation scenes were selected for subsequent rounds of testing and evaluation.

6.1 Testing Round 1

In this first testing round the best individuals from the last 500 generations of eleven successful runs were subject to a uniform set of eight road shapes in each of the twelve scenes. The roads were generated to be approximately 24 m long. For each individual the inherent contrast levels in the scenes were kept the same as that they had experienced during evolution. The response of controllers to previously unseen lower contrast levels is discussed later in Sect. 6.2. The road shapes consisted of two basic types, an “S” shaped course where the robot needed to make turns in both directions to reach the end and the other where there was a constant turn in one direction followed by a straightening of the path. Each of these was generated twice with initial left and right turns for two different angles (20° and 30°) which dictated the curvature of these turns. During evolution the angle of curvature was always 30° and the road generation algorithm ensured that the overwhelming majority (6 out of 7) of shapes generated would be of the first “S” shaped type. The rationale behind generating this fixed set of road shapes was to discover the actual best performing individuals in the population. It was possible that some of the individuals which had obtained high fitness values could have simply been lucky and not possessed the ability to navigate multiple road shapes across all the environments. The re-evaluation tests also provided data on the performance of individuals in each of the twelve scenes, which gave an insight on the effectiveness and flexibility of their dynamic colour perception strategies.

A normalized distance score ranging from 0 to 10 was used to assess performance in each testing condition. Individuals that managed to reach the end of the road in a particular scene would thus get the highest possible score of 10. Figure 3 shows the average of this normalized distance score in each of the twelve scenes. Only data for solutions of evolutionary runs that used six scenes is included here. Figure 4 shows the same, but for solutions when twelve scenes were used during evolution. As during the evolutionary stage, the number of time steps (iterations) in each trial was fixed at 250. Thus individuals with higher scores not only demonstrated better strategies to stay within the road-boundaries but also greater speeds as they moved along the course.

Three out of the ten evolutionary runs using only six-scenes, provided solutions which could solve the three basic mono-colour scenes (road brighter than non-road) and all six dual-colour scenes (Fig. 3). This included scenes 10, 11 and 12 which they had not experienced during evolution. This is proof of the flexibility and adaptability of the solutions evolved. Not surprisingly they failed in the three reversed mono-colour scenes as the entire basis of their learning was dependent on the road being brighter than the non-road. On investigating the dynamic colour perception strategies of these controllers it was observed that the colour outputs for the three mono-colour scenes were more or less steady and above 0.85 throughout the trials. This was expected given their Δ values (see Sect. 5) from evolution being in the range of 1.4–1.8. However in places where sharp turns or course corrections were needed, a different behaviour




Fig. 3. Distance scores in the first round of testing for all twelve scenes. Shown in this graph are scores of the solutions of the three successful evolutionary runs using six scenes.




Fig. 4. Distance scores in the first round of testing for all twelve scenes. Shown in this graph are scores of the solutions of the four best evolutionary runs using twelve scenes.

was observed. The colour nodes instead of staying at a constant high value, oscillated between 0–0.9 every two time steps.

The inclusion of the mono colour scenes and the colour term (Δ) ensured that an adaptive strategy with utilization of all three colour output nodes was developed. It could be argued that including Δ in the fitness function was in a sense dictating a solution to the controllers, rather than truly allowing them to evolve their own strategy. However as seen from the results and during evolution it was indeed a necessary inclusion. Moreover the network did not completely adopt this enforced strategy as suggested by the presence of the periods of oscillatory behaviour displayed by the three colour outputs. It is interesting to note that the motion in terms of dynamics was smoother and faster when the correct colour output was constantly at a high value (≈ 0.9). During the oscillating phases the motion was slower and more uneven, with regular course-corrections having to be made.

The results of the twelve-scene experiments (Fig. 4) were not as uniform, with solutions showing greater variability in their colour perception strategies, depending on the seed and colour distribution set they were evolved in. The majority of solutions (like S5 and S6) only evolved the ability to dynamically vary two of their three colour outputs and simply did not use the third. This meant that two out of the six mono-colour scenes (basic and reversed) could not be solved. The unused colour output varied with each solution. However they did manage to solve all six dual-colour scenes because having the ability to dynamically vary only two colour outputs would be sufficient for these cases.

Only two solutions successfully evolved to show capability of solving all twelve scenes. Of these S4 evolved in scenes with colour distribution of Set *B* and S7 with distribution values of Set *A*. It is interesting to note the effect of these distribution values on the evolved solutions. The seed for S4 when used to evolve a solution with contrast values of Set *A* could develop only a sub-par solution where the controllers could not navigate the green mono-colour scenes. The seed for S7 when used with Set *B*, which could be said to be a less challenging environment, could only solve two scenes. Also unsurprisingly none of these seeds when tried with Set *C* could produce any solutions, as the contrast values were much lower and the distributions themselves were more spread out across the intensity spectrum.

Solution S7, developed a strategy wherein their ability to differentiate on the basis of the green channel was more enhanced than the other two channels. The β output was constant and near maximum for all scenes where bright green could be made the differentiating channel. For all other scenes, the colour outputs oscillated between high and low activations every third time step. While the controllers did traverse the entire course in scenes 1 and 2, the navigation was slower and often error-prone at the beginning, contributing to the lower average scores. Solution S4 evolved behaviour where the α , β and γ terms were near maximum for the majority of the time for scenes 3, 8 and 9 respectively. In the rest of the scenes it displayed periods of both stable and oscillatory activations of the colour output nodes.

6.2 Testing Round 2

Four individuals, two each from the two best six-scene and twelve-scene runs, were then chosen to be subject to a further round of testing. The aim of this round was to investigate the robustness and generality of their road-following strategies by observing their behaviour in environments they had not encountered during the evolutionary phase. The twelve scenes were recreated with textures having average contrast of 90 and deviations from mean of around 40 (on a scale of 0–255). In each of these scenes, the range of distribution of the random noise channels was set at 0–0.80 for one case and 0–0.25 in another. In the evolutionary runs, the distribution of the random noise channels always varied from 0–1 with uniform probability. However it was observed that narrowing this range to 0–0.5 during the testing phase caused a few randomly selected controllers to fail and thus it was decided to add this as a further evaluation parameter. In




Fig. 5. Average scores received by the two best solutions in the second round of testing. The error bars depict the associated standard deviation values.

theory, controllers with the correct feature extraction strategy would be able to completely discard the random channels, as despite the range of the distribution it had no contribution towards highlighting the desired features. The road was set to be of the “S” shaped type with an angle of curvature of 25° in both left and right initial starting directions. These shapes were generated twice, giving a total of 4 trials for each random noise distribution value in each of the 12 scenes. Thus each individual in this second round of testing was evaluated for 96 trials. In order to further enhance the effect of presenting an unfamiliar environment to the controllers the road tile used in this testing phase represented a more delineated and unstructured course, having a maximum width of 110 cm at places but with only 85–90 cm consistently visible throughout. Figure 5 shows the distance scores of the two best solutions of this round, averaged across eight trials for each scene.

The results of this second round of testing (Fig. 5) showed that solutions S4 (twelve scenes) and S7 (twelve scenes) had developed the most robust and general-purpose solution. Despite receiving lower scores (below 7) for a few scenes, only these solutions had the capability of solving all twelve scenes across all the evaluation parameters, i.e. all road shapes with reduced contrast and varying random noise values. The performance of S1 (six scenes) in identifying features in the blue or γ channel was affected by the reduced contrast in the colour distribution. This in turn not only meant failure in the corresponding mono-colour scenes but also in the two dual-colour scenes where the blue channel was brighter on the road. The other two channels could still be successfully used across both ranges of the random noise variation. It was later tested in a scene with average contrast for the blue channel at 109 (still a new environment), and in this case it was able to navigate the corresponding scenes successfully.

While the solution S2 was able to solve almost all scenes when the random noise was in the range of 0–0.80, it failed to differentiate on the basis of both blue and green channels when this range was reduced to 0–0.25. This resulted in lower average scores for scenes 1, 2, 6 and 12. The inability to perform in

these scenes was because it incorrectly associated the low distribution range of random values in the red channel with the availability of features. Thus it could see no contrast between the road and non-road surfaces in those scenes where red was not a feature differentiating channel.

For the two successful solutions in this round, it can be seen that in both cases performance in all but one mono-colour scene deteriorated compared to the earlier round of testing. They were still capable of reaching the end of the road in these scenes, but with less consistency compared to the earlier tests contributing to the lower overall score. Interestingly despite being subject to higher contrasts than S7 during evolution, S4 was still able to match or exceed its performance in eleven out of twelve scenes. On observing the behaviour of the controllers in these lower contrast scenes, it was seen that there was a disparity in their sensitivity to the three colour channels. For each solution there was a particular colour channel in which the ability to perceive contrast was much more pronounced. The controllers changed their colour perception strategies in these scenes, relying increasingly on oscillating the activations of the colour output neurons. However when the channel they were most sensitive to was available, they used it exclusively by activating only the associated output neuron for the majority of the trial.

7 Conclusions

The methodology described in this paper was successful in evolving neural networks capable of demonstrating road-following by dynamic dimensionality reduction in a variety of challenging simulated environments. This new set of controllers have shown improvement in the dynamic colour perceptions abilities compared to those evolved earlier in [1], with the capability to now recognize features based on negative and positive contrast in all three primary colour channels used. These improved results were brought about by the careful design of simulation scenes as well as the formulation of the new fitness function incorporating the colour-term Δ . This work is a significant step towards the hardware implementation of these controllers, as real-world environments would in majority consist of colour combinations similar to those present in the simulated scenes. However it is acknowledged that the contrasts between road and non-road surfaces would be lower than what the networks were tested on. This is proposed to be mitigated by introducing a simple contrast stretching step before the processing of the inputs. Future work would also focus on representing the environment in terms of alternate colour models such as HSV, instead of the traditional RGB model used thus far. Besides this, there is a need to increase the robustness of these controllers by minimizing the disparity in the feature extraction capabilities across the three channels. On the whole however, the findings of this paper strengthen the potential of using these controllers as a viable alternative road-following solution and further efforts would focus on transferring these evolved controllers to a mobile robotic platform.

References

1. Clarke, S., Labrosse, F., Trianni, V., Tuci, E.: An evolutionary approach to road following: a simulated case study. In: 12th European Conference on Artificial Life, pp. 1017–1024. MIT Press, Taormina (2013)
2. Pomerleau, D.: Neural network vision for robot driving. In: Hebert, M.H., Thorpe, C., Stentz, A. (eds.) *Intelligent Unmanned Ground Vehicles*, vol. 388, pp. 1–22. Springer, US (1997)
3. Sotelo, M., Rodriguez, F., Magdalena, L., Bergassa, L., Boquete, L.: A color vision-based lane tracking system for autonomous driving on unmarked roads. *Auton. Robots* **16**(1), 95–116 (2004)
4. Clarke, S.: Thesis: An Evolutionary Approach to Road Following. Aberystwyth University, Aberystwyth (2012)
5. Ososinski, M., Labrosse, F.: Real-time autonomous colour-based following of ill-defined roads. In: Herrmann, G., Studley, M., Pearson, M., Conn, A., Melhuish, C., Witkowski, M., Kim, J.-H., Vadakkepat, P. (eds.) *TAROS-FIRA 2012. LNCS*, vol. 7429, pp. 366–376. Springer, Heidelberg (2012)
6. Ramström, O., Christensen, H.: A method for following of unmarked roads. In: *IEEE Intelligent Vehicles Symposium Proceedings. IEEE* (2005)
7. Pomerleau, D.A., Thorpe, C.E.: Vision-based neural network road and intersection detection and traversal. In: *Proceedings of the International Conference on Intelligent Robots and Systems*, vol. 3, pp. 3344. IEEE (1995)
8. Montemerlo, M., Thrun, S., Dahlkamp, H., Stavens, D.: Winning the DARPA grand challenge with an AI robot. In: *Proceedings of the AAAI National Conference on Artificial Intelligence*, pp. 17–20. AAAI Press (2006)
9. Dudek, G., Jenkin, M.: *Computational Principles of Mobile Robotics*. Cambridge University Press, Cambridge (2000). ISBN 0521568765

Artificial Life and Intelligent Agents

First International Symposium, ALIA 2014, Bangor, UK,

November 5-6, 2014. Revised Selected Papers

Headleand, C.J.; Teahan, W.; Ap Cenydd, L. (Eds.)

2015, XI, 141 p. 55 illus., Softcover

ISBN: 978-3-319-18083-0