Evolving Collective Driving Behaviors

(Extended Abstract)

Chien-Lun (Allen) Huang, Geoff Nitschke
Department of Computer Science
University of Cape Town
Cape Town, South Africa
allen@allenhuang.net, gnitschke@cs.uct.ac.za

ABSTRACT

Recently there has been increased research interest in developing autonomous, adaptive control systems of self-driving vehicles. However, there has been little work on synthesizing collective behaviours for autonomous vehicles that must safely interact and coordinate so as traffic throughput on any given road network is maximized. This work uses neuro-evolution to automate car controller design, testing various vehicle sensor configurations and collective driving behaviours resulting from car interactions on roads without constraints of traffic lights, stop signals at intersections or lanes that vehicles must adhere to and thus simulates potential future scenarios where vehicles must drive autonomously without special road infrastructure constraints. Results indicate that neuro-evolution is an effective method for automatically synthesizing collective driving behaviours that are behaviourally robust across a range of vehicle sensor configurations and generalize to different task environments.

Categories and Subject Descriptors

I.2.6 [**Computing Methodologies**]: Artificial Intelligence, Distributed Artificial Intelligence, Cooperation and Coordination

Keywords

Neuro-Evolution, Self-Driving Vehicles, Multi-Agent Systems

Neuro-Evolution (NE) for Collective Driving

Recent research has focused on producing heuristic adaptive control systems for autonomous vehicles that coordinate their interactions in order to avoid collisions and safely navigate intersections without traffic lights or stop signs [4].

Another approach is to automate the synthesis of vehicle controllers so when vehicles interact a desired collective behavior emerges for any given (road) environment. We present the synthesis of collective driving behaviours using neuro-evolution. Task performance is measured by the total vehicle throughput between transit points on given roads. We also investigate the impact of vehicle *morphological complexity* [1] (sensory configuration) on task performance. *Neuro-Evolution for Augmenting Topologies* (NEAT) [5] is used for controller evolution as it has been applied in similar studies [2]. Controllers were evolved to maximise the average distance traversed on tracks with obstacles and oncoming traffic, whilst minimising collisions. Sensory input and motor output layers were fixed during evolution and NEAT adapted the number of hidden layer nodes and connectivity between inputs and outputs (figure 1). Controller evolution used one of four pre-determined sensor configurations (figure 1) each increasing in number of sensors (morphological complexity). Vehicles used heuristic driving behaviours (via check-points) but if imminent collisions were detected NEAT evolved behavior overrode the heuristics.

Appears in: *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), Das, Durfee, Larson, Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

Experiments

An extension of *UnityNEAT* based on *SharpNEAT* was used to simulate physically realistic 3D vehicles, sensors and roads. Homogenous vehicle groups (each vehicle had the same controller and morphology) were evolved to avoid collisions with obstacles, road-side barriers and other vehicles, while traversing the road as quickly as possible. Controllers were awarded fitness equalling the portion covered of the track's length (via check-points) over 150 simulation (task trial) iterations (equation 1), where cp_{passed} is the number of check-points vehicles successfully pass, cp_{total} is the total number of check-points and *coll* is the number of collisions.

$$\text{Fitness}(x) = \frac{1}{cars} \sum_{i=0}^{cars} \left(\frac{cp_{passed}}{cp_{total}} * 0.9^{coll}\right) \quad (1)$$

One generation comprised the evaluation of all controllers on a track and there was one simulation task trial (for each controller) per generation. Each generation, 12 vehicles were initialized at given starting points on a track. Four sensory configurations (morphologies) were tested, ranging from simple to complex (figure 1). Each sensor was a simulated radar with a pyramidal sensory *Field of View* (FOV), detecting the closest object to the vehicle. The symmetrical sensory configuration was selected given similar designs in previous research [2].

Two tracks representing realistic driving environments were used in these experiments. Track 1 was a straight road of varying width and elevation, with two starting points where six vehicles initialised at each end. Track 2 was a four-way intersection where vehicles had to converge from four starting points at differing elevations. Three vehicle were initialized at each starting point and followed different routes to various destinations so vehicles had to coordinate to safely

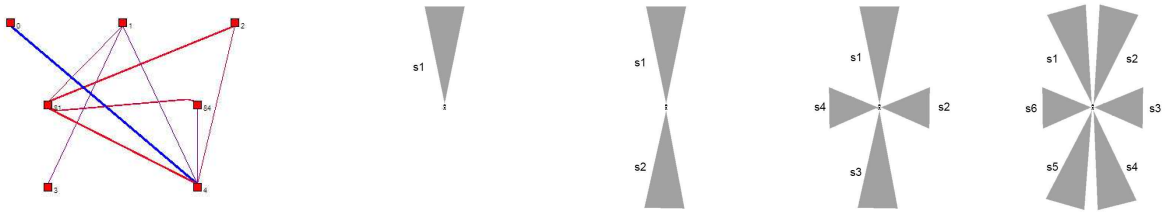

Figure 1: ***LEFT: Example Evolved Controller:*** **Top three nodes represent *bias* and *sensors [1, 2]* inputs respectively. Bottom nodes represent *steering* and *Braking* overrides. NEAT adapts connections (links and weights) and hidden layers.** ***RIGHT: Vehicle Sensory Configurations.*** **From left to right, numbers of sensors increase and thus increase controller input node complexity. Sensors facing the direction of travel have a $100m$ range and $40°$ FOV, whilst side sensors have a $50m$ range a FOV of $20°$.**

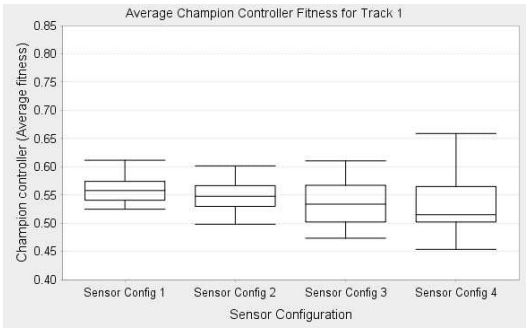

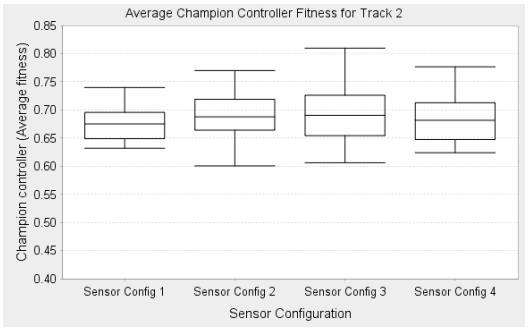

Figure 2: Average task performance (normalized in the range: [0, 1]) of collective driving behaviors evolved for each sensory configuration on track 1 (*left*) and track 2 (*right*). On track 1, *Mann-Whitney U*, $p \leq 0.05$ statistical tests indicated statistical difference between configurations [1, 2] and [3, 4]. On track 2, *Mann-Whitney U*, $p \leq 0.05$ statistical tests indicated no statistical difference between all configurations.

pass through the intersection [4]. On both tracks [1], obstacles appeared at preset points at a distance of approximately half of a vehicle's sensor range. This was to simulate the sudden crossing of pedestrians.

Results and Discussion

Task performance results support this study's first objective of demonstrating the efficacy of NE to synthesize collective self-driving controllers (figure 2). For all sensor configurations on both tracks, evolved collective driving behaviour achieved above median task performance. To address the second objective (to ascertain the impact of *morphological complexity* on evolved collective driving), we gauged relative task performances of behaviours evolved for each sensory configuration. Pair-wise statistical tests (Mann-Whitney, $p \leq 0.05$, [3]) indicated a significant difference in task performance between configurations $[1, 2]$ and $[3, 4]$. Also, to test how well evolved controllers generalized to other tracks, we evaluated *track 1 evolved* controllers on *track 2* and vice-versa. Results indicated that for sensor configurations $[1, 2, 3]$, controllers evolved on both tracks traversed either track with comparable task performance. Future research will evolve controllers on one track and evaluating it on multiple *new* tracks to more rigorously test an evolved controller's capability to generalize to new task environments.

[1] `https://people.cs.uct.ac.za/~chuang/aamas2017`

Acknowledgements

This research was partially funded by the *Centre for Artificial Intelligence Research* (CAIR) initiative under the *Council for Scientific and Industrial Research & Department of Science and Technology of South Africa* (CSIR-DST) Inter-Bursary programme.

REFERENCES

[1] J. Auerbach and J. Bongard. Environmental influence on the evolution of morphological complexity in machines. *PLoS Computational Biology*, 10(1):e1003399. doi:10.1371/journal.pcbi.1003399, 2014.

[2] L. Cardamone, D. Loiacono, and P. Lanzi. Learning to drive in the open racing car simulator using online neuroevolution. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(3):176–190, 2010.

[3] B. Flannery, S. Teukolsky, and W. Vetterling. *Numerical Recipes*. Cambridge University Press, Cambridge, UK, 1986.

[4] M. Middlesworth, K. Dresner, and P. Stone. Replacing the stop sign: Unmanaged intersection control for autonomous vehicles. In *Proceedings of the Workshop on Agents in Traffic and Transportation Multiagent Systems*, pages 94–101, Estoril, Portugal, 2008.

[5] K. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127, 2002.