PERFORMANCE OF POPULATION OF NAÏVE CREATURES WITH FEAR AND DESIRE CAPABLE OF OBSERVATIONAL SOCIAL LEARNING*

Anna T. Lawniczak[a], Bruno N. Di Stefano[b], Leslie Ly[a]
Shengkun Xie[c]

[a]Department of Mathematics and Statistics, University of Guelph
Guelph, Ontario, Canada
[b]Nuptek Systems Ltd., Toronto, Ontario, Canada
[c]Ted Rogers Management School, Ryerson University
Toronto, Ontario, Canada

(Received March 1, 2016)

In the microscopic modeling of swarm of robots, individual robots may be identified as cognitive agents. We describe a model of a population of simple cognitive agents, naïve creatures, experiencing fear and/or desire when learning to cross a highway. The creatures use an *observational social learning* mechanism in their decision to cross the highway or not. We study in various traffic environments characterized by vehicle traffic density, how fear and/or desire affects creatures' performance measured by the number of killed, queued and successful creatures. We investigate how this performance is affected when the creatures' knowledge accumulated in one traffic environment is transferred to a different traffic environment for them to continue their learning. We consider the case when creatures are not allowed to change their crossing point and when they are allowed to change it.

DOI:10.5506/APhysPolBSupp.9.95

1. Introduction

Research in swarm robotics has shown that, for carrying out some tasks (*e.g.*, target or source search, task allocation, exploration, mapping, co-operative transportation, unmanned aerial vehicle, *i.e.* UAV, controlling, post-disaster relief), it may be more efficient, reliable and economical to employ a large number (hundreds or thousands) of very simple robots than

* Presented at the Summer Solstice 2015 International Conference on Discrete Models of Complex Systems, Toronto, Ontario, Canada, June 17–19, 2015.

to employ a small number of sophisticated ones, see [1, 2] and references therein. Often, an autonomous swarm of robots have to operate in unknown environments to which they have to adapt through learning. Thus, learning plays an important role in swarm robotics. Since individual robots in a swarm are usually architecturally minimal with limited computational capabilities, it is important that, in a swarm of robots, the implemented learning algorithms are not computationally demanding. In the microscopic modeling of swarm of robots, individual robots may be identified as cognitive agents capable of performing cognitive acts; *i.e.* a sequence of the following activities: (1) *Perceiving* information originating from the environment and from other agents; (2) *Reasoning* about this information using existing knowledge; (3) *Judging* the obtained information using existing knowledge; (4) *Responding* to other cognitive agents or to the external environment, as it may be required; (5) *Learning*; *i.e.* changing (and hopefully augmenting) the existing knowledge if the newly acquired information allows it [3–6]. The functionality and performance requirements of cognitive agents were studied in [7]. The architecture of an application in the independent software implementation of a generic cognitive agent able of providing the required functionality and performance was proposed in [8].

The aim of this work is to study how the performance of a simple learning algorithm based on an *observational social learning* mechanism [9–11], used by cognitive agents in the model developed in [12–16] is affected by agents' fear and/or desire to cross a highway. The performance of the agents, called "naïve creatures", is measured by numbers of killed, queued and successful creatures. We investigate the creatures' performance under various traffic conditions, characterized by car density, and we study how this performance is affected when the creatures' knowledge accumulated in one traffic environment is transferred to a different traffic environment for them to continue their learning. We consider the case when creatures are not allowed to change their crossing point and when they are allowed to change it.

The paper is structured as follows: Section 2 describes briefly the considered model of naïve creatures learning to cross a highway, their learning algorithm and their virtual environment; Section 3 presents selected simulation results; Section 4 provides our conclusions and outlines the future work.

2. Model of naïve creatures learning to cross a highway

We review here only the main features of the model reported in [12–16] and refer the reader for details there.

In this work, we assume that the environment is a single-lane unidirectional highway without intersections. We model the highway traffic by adopting the Nagel–Schreckenberg cellular automaton model and refer the reader to [17–20] for the details. The model consists of four steps that are applied simultaneously to all cars: acceleration, safety distance adjustment, randomization, and change of position. The cars are generated randomly at "starting cells" with car creation probability (Car Prob.) and are assigned a random speed between zero and the maximum allowed speed for cars, which is set in the configuration file.

At each time step, a creature is generated only at the crossing point (CP) set at the initialization step and is placed into the queue at this CP. Each generated creature falls with equal probability (0.25) into one of the four categories: (1) no fear nor desire; (2) only fear; (3) only desire; (4) both fear and desire. The creatures' attributes play a role in their decision making process on whether or not to cross the highway through the values of fear (aversion to risk taking) and desire (propensity to risk taking) that creature may experience. Creatures want to cross the highway without being killed by the oncoming vehicles and they have a strong instinct to survive.

Each creature is *an autonomous entity capable of interacting with its environment and other agents* capable of: (1) matching simple patterns; (2) evaluating distances in an approximate way; (3) evaluating the velocity of moving vehicles in an approximate way; (4) assigning a discrete number to an approximate class; (5) witnessing what happened to the creatures that previously crossed the highway at this crossing point (with exclusion of the first creature); (6) evaluating what they witnessed in (5), *i.e.* if it was successful or not; (7) imitating the creatures which crossed successfully; (8) deciding not to cross and wait for better conditions or to look for a different crossing point when unsuccessful crossings outnumbered the successful ones. All of the above allow each crossing point (CP) to build one knowledge base (KB) during the experiment that is available to all creatures at that CP.

The creatures attempt to cross the highway having a limited horizon of vision and perceiving only fuzzy levels of distance (*e.g.*, "close", "medium", "far") of cars within this horizon and their speeds (*e.g.*, "slow", "medium", "fast"). The ranges for these qualitative categories are set in the simulator's configuration file. The creatures may build up in the queue as a result of not crossing at each time step. If the simulation setup permits, after deciding not to cross the highway, a creature may move randomly along the highway horizontally in either direction to a new CP or it may stay at the same CP with equal probability of 1/3. The number of horizontal cells a creature may move in one-time step is 1 and the maximum distance the creature may deviate from its original CP in both directions is 5. If the creature at the top of a queue leaves the queue, the creature that was behind moves to the top of the queue. When a creature crosses the highway at a given CP, information is recorded into the knowledge base (KB) of all the creatures at this CP. The information about qualitative description of velocity (*e.g.*, such as "fast", "medium" and "slow") and of the distance (*e.g.*, such as "close", "medium", "far") is stored, respectively, in the columns and rows of the KB table.

The KB table is initialized as *tabula rasa*, *i.e.* with all its entries set to 0, allowing creatures to cross the highway regardless of the observed (distance, velocity) levels until the first successful crossing of a creature, or five consecutive unsuccessful crossings, whichever comes first. If a creature successfully crossed the highway, the perceived (distance, velocity) score in the KB table is increased by one point. If the creature was killed, it is decreased by one point.

After the initialization of the simulation, each creature at the top of the queue consults the KB table to decide if it is safe or not to cross. Its decision is based on the implemented intelligence/decision making algorithm, which for a given (distance, velocity) pair combines the *success ratio* of crossing the highway for this (distance, velocity) pair with the creature's fear and/or desire values, as follows.

For each (distance, velocity) pair at each time step, the numerator in the success ratio is the current value from the KB table, *i.e.* it is the number of *successful crossing* minus the number of *unsuccessful crossings* for this (distance, velocity) pair up to this time. The denominator is the total number of creatures that have crossed the highway successfully, regardless of the (distance, velocity) combination, up to this time; *i.e.* it is the number describing the creatures' entire population success up to this time. If for some (distance, velocity) configuration at the simulation start, all creatures are killed, ratio becomes "$-5/0$". In this case, we set the success ratio to zero since "division by zero" is undefined.

A randomly generated creature will base its decision on the formula: (1) *success ratio* + *value of desire* − *value of fear*, if it has both fear and desire; (2) *success ratio* − *value of fear*, if it has only fear; (3) *success ratio* + *value of desire*, if it has only desire; (4) *success ratio*, if it has no fear and no desire. If a creature and a given (distance, velocity) combination yield from the formula a value that is less than zero, then the creature will not attempt to cross the highway under this condition and it will wait for a configuration for which the value of the formula is non-negative, or it may decide to move to another crossing point.

The main simulation loop of the model consists of: (1) generation of cars at each lane of the highway using the Car Prob.; (2) generation of creatures at each CP with their attributes; (3) update of the car speeds according the Nagel–Schreckenberg model; (4) movement of the creatures from the CP queues into the highway (if the decision algorithm indicates this should occur); (5) update of locations of the cars on the highway. This includes passing other cars in the case of multi-lane highway (which is not considered here). It also includes the logic to check if any creature has been killed; (6) advancement of the current time step. After the simulation has been completed, the results are written to output files using an output function.

3. Learning performance of naïve creatures

We present selected simulation results showing the learning performance of a population of naïve creatures experiencing various levels of fear and/or desire when learning to cross a single-lane unidirectional traffic highway. The creatures' performance is measured by the average number of killed, queued and successful creatures at simulation end. We run each simulation for 1511 time steps and repeat it 30 times for different random seed values. We investigate the creatures' performance under various traffic conditions, characterized by car density, and we study how this performance is affected by the transfer of the creatures' knowledge accumulated in one traffic environment to a different traffic environment where they will continue to learn. We consider the case when creatures are not allowed to change their crossing point and when they are allowed to change it.

The highway has a length of 120 cells and as customary in traffic engineering literature, [17], each cell represents a segment of a highway of 7.5 m in length. Thus, the highway is 900 m long. A car is generated at the start of the highway at each time step with a given Car Prob., which varies between the values: 0.1, 0.3, 0.5, 0.7, and 0.9. The Car Prob. determines the highway car traffic density. The maximum speed that a car may achieve is 11 cells per time step. We consider the traffic without erratic drivers.

Here, we present simulation result only for one CP setup at cell 80 at the initialization step, *i.e.* 600 m away from the beginning of the highway. The distance from where the vehicles are generated is important because it will affect the nature of the vehicle traffic. For example, there will likely be more vehicles traveling at maximum speed and in a more homogeneous manner near CP 80 than at CP 40.

The value of fear and value of desire parameters both vary between the values: 0, 0.25, 0.5, 0.75, and 1. Being a part of the decision formula, these values influence the creatures' decision making process of whether or not to cross.

The KB table is 3 by 4 with an extra entry, as it has 3 groupings of distance and 4 groupings of speed. The creatures in this case specifically perceive: (1) "close" for a vehicle being 0 to 3 cells away, "medium" if it is 4 to 5 cells away, "far" if it is 6 to 7 cells away and "out of range" (the extra entry) if it is 8 or more cells away, regardless of the velocity of the vehicle; (2) "slow" when the perceived velocity of a vehicle is 0 to 3 cells per time step, "medium" when it is 4 or 5 cells per time step, "fast" when it is 6 or 7 cells per time step, and "very fast" 8 to 11 (maximum speed) cells per time step.

We consider here the case in which, for each given value of Car Prob., the KB tables are always transferred from the current repetition to the next one within any particular configuration of the parameters' values. We call this Framework (II), to distinguish it from Framework (I) where such transfer does not happen. Results of Framework (I) are discussed elsewhere. Thus, in Framework (II) the n^{th} repeated simulation starts with the KB table accumulated over "$n - 1$" previous simulation runs. The distinction between Framework (I) and Framework (II) is important as the amount of learning under each Framework is different and this affects the creatures' success in crossing the highway. In Framework (II), the KB tables become much more developed as there is much more transferring of the KB tables occurring. In Framework (II), for each given value of Car Prob., the KB tables are built by several populations of creatures, depending on a number of considered repeats. If at the end of all the considered repeats for a given value of Car Prob., the KB table is not transferred to the beginning of the simulation at a different traffic environment (*i.e.*, with a different value of Car Prob.), we say that KB transfer is "none". If in Framework (II) the KB table is additionally transferred at the end of all the repetitions in the environment with lower (higher) Car Prob. to the one with an immediately higher (immediately lower) Car Prob., we say that KB transfer is "forward" ("backward").

Here, we discussed selected simulation results obtained under Framework (II) for the cases of KB transfer "none" and "forward". The results for KB transfer "backward", comparison of creatures' performance under the two frameworks and more detailed analysis how various model parameters effect creatures' performance will be discussed elsewhere.

The enclosed figures display respectively in the columns of each figure, plots of the average numbers of killed creatures (first column), queued creatures (second column), successful creatures (third column) and one standard deviation above the mean (displayed by black colour bars attached to the other colour bars). The figures' rows display respectively, the plots for different values of Car Prob. The values of Car Prob. are listed directly under each plot. Each plot displays the mean values and one standard deviation above the mean as a function of the considered fear and desire values. There are 25 mean values and one standard deviations displayed on each plot.

Figure 1 displays the results when KB transfer is "none" and creatures are not allowed to move horizontally. Figure 2 displays the results when KB transfer is "forward" and creatures are not allowed to move horizontally. Figure 3 displays the results when KB transfer is "none" and creatures are allowed to move horizontally. Figure 4 displays the results when KB transfer is "forward" and creatures are allowed to move horizontally.

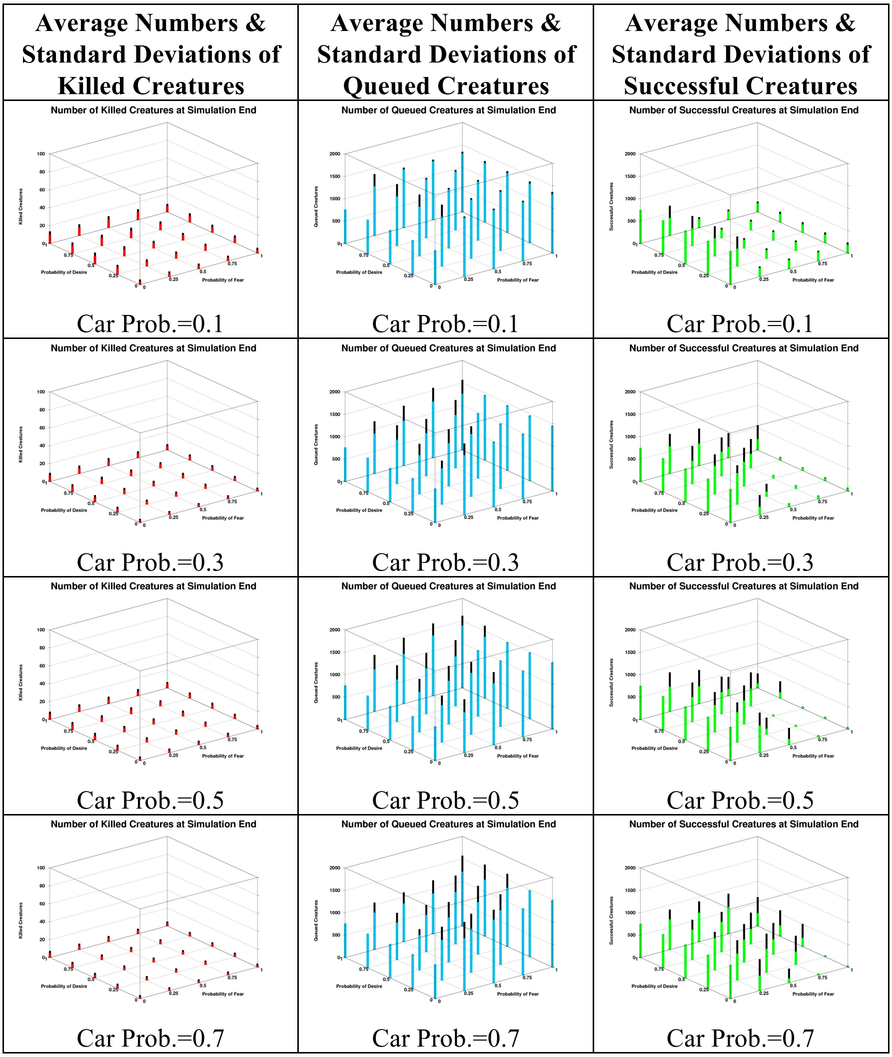

Fig. 1. Plots of average numbers and one standard deviations (in black) of killed, queued and successful creatures, respectively in columns, as function of creatures' fear and desire values, at the end of simulation for single-lane 120 cells-long highway without erratic drivers, at the crossing point cell 80. Creatures are not allowed to move horizontally. KB table is not transferred "forward" from one environment to the next one characterized by car creation probability listed under each plot. The statistics were calculated over 30 simulations.

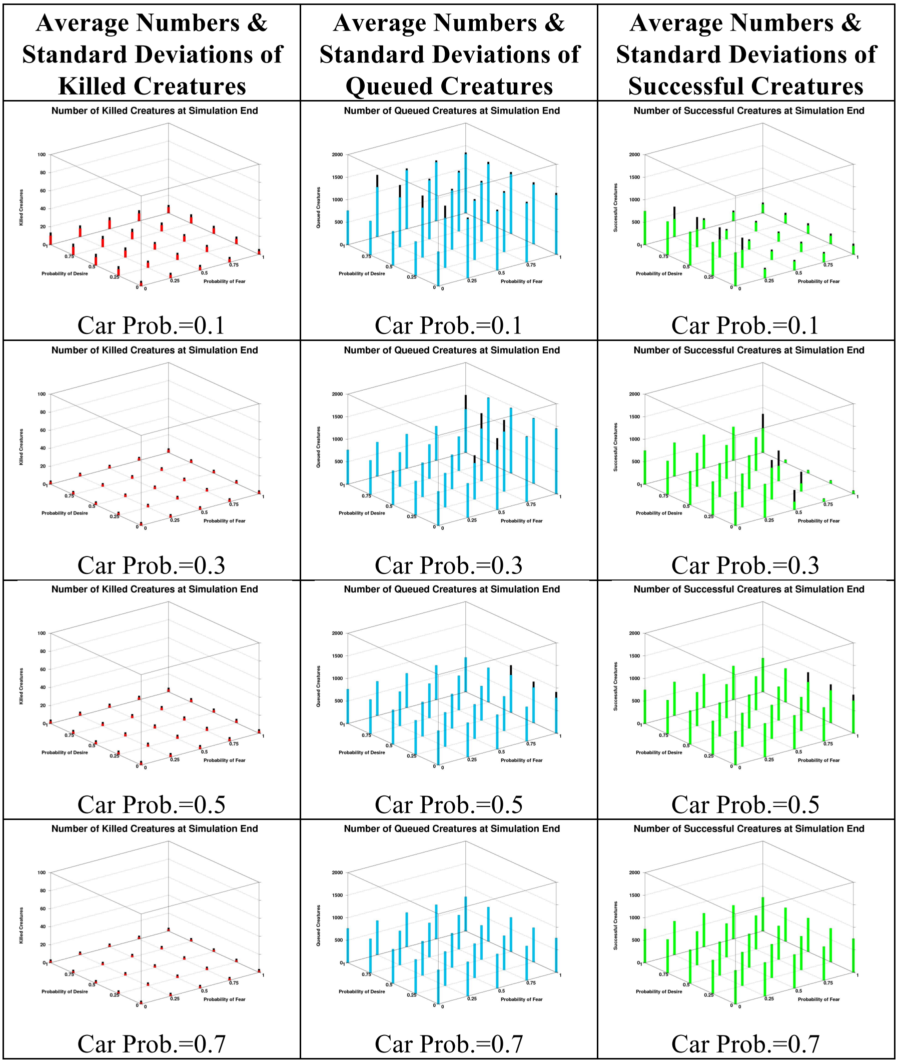

Fig. 2. Plots of average numbers and one standard deviations (in black) of killed, queued and successful creatures, respectively in columns, as function of creatures' fear and desire values at the end of simulation for single-lane 120 cells-long highway without erratic drivers, averaged over 30 runs, at the crossing point cell 80. Creatures are not allowed to move horizontally. KB table is transferred "forward" from one environment to the next one characterized by car creation probability listed under each plot. The statistics were calculated over 30 simulations.

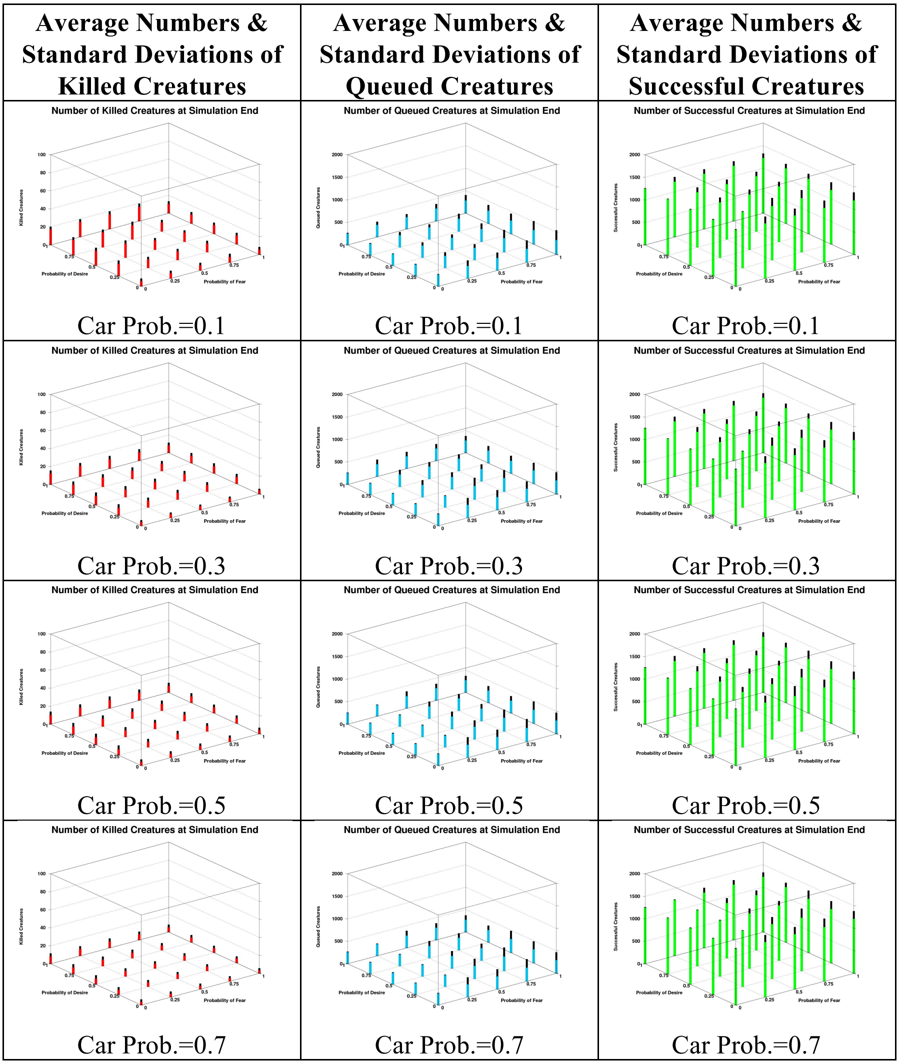

Fig. 3. Plots of average numbers and one standard deviations (in black) of killed, queued and successful creatures, respectively in columns, as function of creatures' fear and desire values at the end of simulation for single-lane 120 cells-long highway without erratic drivers, at the crossing point cell 80. Creatures are allowed to move horizontally to a different crossing point. KB table is not transferred "forward" from one environment to the next one characterized by car creation probability listed under each plot. The statistics were calculated over 30 simulations.

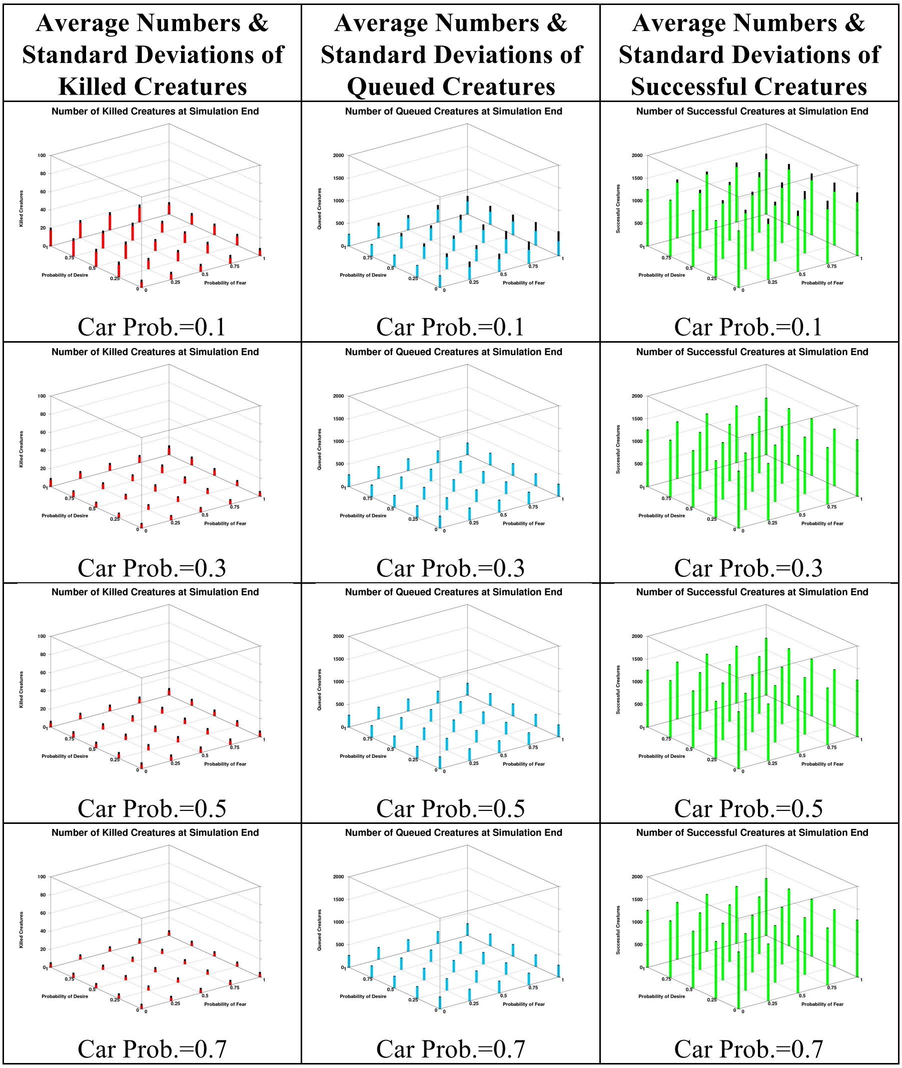

Fig. 4. Plots of average numbers and one standard deviations (in black) of killed, queued and successful creatures, respectively in columns, as function of creatures' fear and desire values at the end of simulation for single-lane 120 cells-long highway without erratic drivers, at the crossing point cell 80. Creatures are allowed to move horizontally to a different crossing point. KB table is transferred "forward" from one environment to the next one characterized by car creation probability listed under each plot. The statistics were calculated over 30 simulations.

The presented simulation results show that for killed creatures: (1) their mean-value numbers and one standard deviations are very small in comparison with mean-value numbers and one standard deviations for the other types of creatures, regardless of fear and desire values and Car Prob. values; (2) the mean-value numbers and one standard deviations increase when creatures are allowed to move to a different crossing point, as more creatures are attempting to cross; (3) these numbers decrease monotonically with increase in Car Prob. values, because more creatures queue to cross when Car Prob. values increase; (4) the decrease in (3) is enhanced by KB transfer "forward", as more creatures are successfully crossing the highway and less of them are queuing due to the additional transferring of knowledge.

The presented in Fig. 1 simulation results (KB transfer "none") show that fear (aversion to risk taking) has very negative effect on the number of creatures successfully crossing the highway, in the presence of fear, the majority of the creatures choose to wait (leading to large queues), regardless of creatures desires (propensity to risk taking). This effect is more dominant for lower values of Car Prob. Higher values of desire improve the creatures' ability to cross successfully and slightly decreases the queues. This improvement becomes much more significant in the case when the KB transfer is "forward", as can be seen in Fig. 2. Regardless of the value of fear, more creatures are crossing successfully instead of being queued as the KB tables become more developed through passing them from the environment with a lower Car Prob. to the one immediately higher. The system reaches sort of uniformity in distribution of successful and queued mean values and their variabilities for Car Prob. 0.7 (*i.e.*, sizes of the bars), which is maintained almost unchanged for Car Prob. 0.9 (results not displayed here).

When creatures are allowed to move to different crossing points, more creatures attempt to cross the highway than when they are not allowed to move. Thus, more of them may cross successfully regardless of their fear values as can be seen from Fig. 3. The sizes of mean values of successful creatures are high and the sizes of queued creatures are rather low regardless of the fear and Car Prob. values. However, there is a visible variability present in the mean values of successful and queued creatures' numbers when the creatures have fear. This variability diminishes with KB transfer "forward" as can be seen from Fig. 4. More KB is transfer "forward" less variability is observed.

We presented here, selected simulation results and focused on KB transfer "forward". In the case of KB transfer "backward", the qualitative effects of KB transfer are similar with their evolution taking place from higher to lower Car Prob. values. In the case of Framework (I), as much less training takes place within the same environment, the KB transfer has less strong positive effect on the system performance.

4. Conclusions and future work

Our simulations show that the presence of fear has a depressing effect on the system performance and cannot be easily compensated by desire in the case when only one crossing point is considered and KB transfer is "none". The "forward" transfer of KB improves the system performance and when coupled with desire, this improvement happens already for lower values of Car Prob., *i.e.* with less creatures training. In the case when creatures are allowed to move, this gives creatures more opportunities to be successful and more training (KB transfer "forward") decreases system variability.

We presented here selected simulation results focusing on the effects of KB transfer "forward". More extensive and detailed results will be reported elsewhere.

The authors acknowledge help of Hao Wu in preparing the plots, the prior work of Jason Ernst and helpful discussions with Fei Yu. A.T.L. acknowledges partial financial support from the NSERC of Canada. B.N.Di S. acknowledges full financial support from Nuptek Systems Ltd.

REFERENCES

[1] Y. Tan, Z.-y. Zheng, *Def. Technol.* **9**, 18 (2013).

[2] I. Navarro, F. Matia, *ISRN Robotics* **2013**, 608164 (2013).

[3] S. Russell, P. Norvig, *Artificial Intelligence, A Modern Approach*, Pearson Education Limited, 2014.

[4] D.L. Poole, A.K. MacKworth, *Artificial Intelligence, Foundations of Computational Agents*, Cambridge University Press, Cambridge, US, 2010.

[5] J. Ferber, *Multi-Agent Systems. An Introduction to Distributed Artificial Intelligence*, Addison Wesley, London, UK, 1999.

[6] A.M. Uhrmacher, D. Weyns, *Multi-Agent Systems Simulation and Applications*, CRC Press, Taylor & Francis Group, Boca Raton, Fl., 2009.

[7] B.N. Di Stefano, A.T. Lawniczak, Proc. 2009 IEEE TIC-STH, Toronto International Conference, 2009, p. 509.

[8] A.T. Lawniczak, B.N. Di Stefano, *Procedia Comput. Sci.* **1**, 2227 (2010).

[9] Ch.L. Nehavin, K. Dautenhahn, *Imitation and Social Learning in Robots, Humans and Animals*, Cambridge University Press, Cambridge, UK, 2007.

[10] W. Hoppitt, K.N. Laland, *Social Learning, An Introduction to Mechanisms, Methods, and Models*, Princeton University Press, Princeton 2013.

[11] E. Alonso *et al.*, *Knowl. Eng. Rev.* **16**, 277 (2001).

[12] B.N. Di Stefano, A.T. Lawniczak, *Acta Phys. Pol. B Proc. Suppl.* **5**, 21 (2012).

[13] A.T. Lawniczak, J.B. Ernst, B.N. Di Stefano, *Lect. Notes Comput. Phys.* **7495**, 425 (2012).

[14] A.T. Lawniczak, J.B. Ernst, B.N. Di Stefano, *Proceedia Comput. Sci.* **18**, 2611 (2013).

[15] A.T. Lawniczak, B.N. Di Stefano, J.B. Ernst, *Lect. Notes Comput. Phys.* **8751**, 688 (2014).

[16] A.T. Lawniczak, B.N. Di Stefano, J.B. Ernst, IEEE Symposium on Intelligent Agents (IA) 2014, Orlando, FL, 09–12 Dec. 2014, p. 30, DOI: 10.1109/IA.2014.7009455.

[17] K. Nagel, M. Schreckenberg, *J. Physique I* **2**, 2221 (1992).

[18] A.T. Lawniczak, B.N. Di Stefano, Proc. of Automata-2008, Theory and Applications of Cellular Automata, Luniver Press, UK, 2008, p. 527.

[19] A.T. Lawniczak, B.N. Di Stefano, *Lect. Notes Comput. Phys.* **6350**, 600 (2010).

[20] A.T. Lawniczak, B.N. Di Stefano, *Acta Phys. Pol. B Proc. Suppl.* **3**, 479 (2010).