MODELLING A SIMPLE ADAPTIVE COGNITIVE AGENT*

Bruno N. Di Stefano†

Nuptek Systems Ltd, Toronto, Ontario M5R 3M6, Canada
b.distefano@ieee.org

Anna T. Lawniczak

Department of Mathematics and Statistics, University of Guelph
50 Stone Road East, MacNaughton Bldg, Guelph, Ontario N1G 2W1, Canada

(Received December 29, 2011)

Agent-based models approximate the behaviour of simple natural and man-made systems. Their performance is limited by the abstractions used to implement them, *i.e.*, finite state machines. "Cognitive Agents" perform "Cognitive Acts" (*i.e.*, Perceiving, Reasoning, Judging, Responding, and Learning) and are the closest to the behaviour of simple biological entities. We present a simple cognitive agent capable of evaluating if a strategy has been applied successfully and capable of applying this strategy again with small changes to a similar but new situation. We describe how a simple agent can be trained to learn how to safely cross a road with "one lane one directional highway" and later, when the situation changes, a road with "two lanes one directional highway" and "two lanes on a bi-directional highway". Future research is outlined.

DOI:10.5506/APhysPolBSupp.5.21
PACS numbers: 89.20.Kk, 89.20.Bb

1. Introduction

Agent-based models accurately approximate the behaviour of simple or simplified natural and man-made systems by mimicking and often bio-mimicking simple entities. Many definitions exist for the term "agent" in this context. Most of them are not terse and not easy to remember [1]. We like the short definition according to which an agent is "an autonomous entity

* Presented at the Summer Solstice 2011 International Conference on Discrete Models of Complex Systems, Turku, Finland, June 6–10, 2011.

† Corresponding author.

capable of interacting with its environment and other agents" [2] and [3]. The word "interacting" is used in an anthropomorphic sense. It means that the agent communicates with and receives stimuli from other agents and the environment and reacts accordingly to its understanding of the situation. The ability of agents to act autonomously is limited to the predefined environment and to the predefined situations to which the agent is expected to respond, because agents can only act in a situation compatible with the way they are designed. In fact, traditionally the behaviour of the agent is provided by means of a finite state machine or a set of finite state machines [2] and [3]. A problem with all finite state machines is that their design, verification, validation, coding, and testing becomes progressively harder when trying to prepare the finite state machine for all possible scenarios beyond a small number. "Cognitive agents" partially solve this problem by performing "Cognitive Acts" (*i.e.*, a sequence consisting of all of the following acts: Perceiving, Reasoning, Judging, Responding, and Learning) as opposed to agents, who perform "Reflexive Acts" (*i.e.*, Perceiving and Responding) [1]. A cognitive agent has the ability to react to unexpected situations and the ability to reason about these unexpected situations in order to react to them. The functionality and performance of cognitive agents requires replacing the finite state machine typical of "non cognitive" agents, *i.e.* "reactive agents", with more complex functional blocks, built using computational intelligence methodology, *i.e.*: fuzzy logic, neural networks, evolutionary computation, and various types of bio-mimicry. In reality, also cognitive agents are implemented by means of software or a mix of hardware and software and are still far from the performance of animals and humans. All animal species, from insects to mammals, so far outperform digital computers in performing cognitive functions [4]. The reason for this gap is that animal species use a mix of analog and digital functions emphasizing *distributed, event-driven, collective, and massively parallel mechanisms, and make extensive use of adaptation, self organization, and learning* [4]. Indeed, all human beings drive cars with a certain amount of dealing with imprecise, incomplete, and vague information. Moreover, most of the time, human beings reason and act, as a consequence of their reasoning, without taking any precise measurement and without performing any precise computation (*e.g.*, we park cars successfully with a high degree of approximation). Human beings unconsciously map what they perceive (*e.g.*, hear, see, *etc.*) to an idealized view of reality. They apply approximate rules to this approximate view of reality and, eventually, produce an approximate result of their inference [3]. Geoffrey Hinton, one of the pioneers in artificial intelligence, in an interview, said that computers *are not nearly as good* (as humans). *Part of it is the hardware, we have many billions of neurons, each of which has thousands of connections. Even now, it is very hard to get computers that have the* *same amount of processing power and particularly the same access to stored knowledge. The brain can access many gigabytes of knowledge in a tiny fraction of a second. Only the biggest supercomputers can do that kind of thing at present* [5].

However, we believe that a good level of successful bio-mimicking can be achieved with a few basic computational intelligence building blocks. In [3] we wrote this could be achieved by defining and implementing a hierarchical software layered model for the generic cognitive agents, a layered model in a style akin to the hierarchical ISO OSI 7 layer model used in data communication [6]. In this paper, we present a simple adaptive cognitive agent capable of evaluating if a strategy has been applied successfully and capable of applying this strategy again with small changes to a similar but new situation. We believe that a simple cognitive agent able to apply a finite collection of algorithms can be evolved to deal with much more complex situations.

2. The environment, the agent, and the experiment

At first, we assume that our environment consists of a segment of a one lane highway characterized by unidirectional vehicular traffic, without any intersection, see Fig. 1. This expressway and its traffic can be represented and studied with a non-periodic boundary conditions Nagel–Schreckenberg-like model, that is, a Cellular Automaton (CA) model based on an extension of Elementary Cellular Automata (ECA) Rule 184 [7]. As customary in the traffic modelling literature, we model the one lane highway as a large number of adjacent cells, with each cell representing a segment of highway of 7.5 m in length, *e.g.* [7] and [8]. Such representation has been chosen because it corresponds to the space occupied by the typical car plus the distance to the preceding car in a situation of dense traffic jam. The traffic jam density

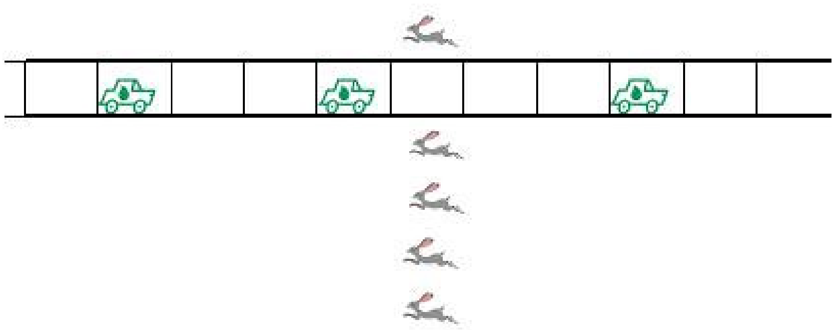

Fig. 1. Simplified diagram showing the highway, several cars, and several agents (*i.e.*, creatures, rabbits).

is given by 1000/7.5 m approximately equal to 133 vehicles per km. The cells can be used to establish a system of 1-D coordinates. We assume that our agent represents a creature with a strong impulse to cross the highway, see Fig. 1. The impulse to cross could depend on some natural instinct or need, *e.g.* foraging, mating, *etc.* There is a population of such creatures. All creatures can cross the highway only at a predefined location defined by a given cell number. Because only one creature can cross the highway at a given moment in time, the creatures must line up in front of the crossing. The instinct to cross the highway may be mitigated by what the creature has witnessed, for instance it may have seen a similar creature being struck by a vehicle while crossing the highway. All creatures lining up to cross the highway witness a certain number of crossings and their outcomes (*i.e.*, safe crossing or death because of colliding with a vehicle). The creatures are not able to measure distances accurately. They can only rank the position of the vehicle with respect to the crossing cell according to a discrete number of categories (*e.g.*, {far, mid range, close} or, alternatively, {very far, far, mid range, close, very close}, *etc.*).

It is reasonable to assume that at the beginning of the simulation many creatures will be struck by vehicles. With the passing of time the creatures learn how to decide if it is safe to cross or if it is better to wait. Eventually, for certain densities of traffic and for certain speeds, it may never be safe to cross.

Our work consists of giving the creatures a very primitive, basic algorithm, to match patterns and to decide when it is safe to cross, based on past experience. The same basic algorithm is expandable to different situations. For instance, if instead of having to cross a one lane highway the creatures are required to cross a two lane highway, the creatures must be able to adapt. Of course, at the transition, that is the change of highway type, a few creatures may be struck by vehicles and perish. However, after a short time they recognize the new pattern and adjust the algorithm. Similarly, if, instead of dealing with a one lane one directional highway, the creatures deal with a multilane bidirectional highway, they can adapt.

2.1. One lane one directional highway

We assume that each creature can observe over time incoming vehicles for a sufficiently long time to be able to infer the approximate speed of these vehicles. At time we call by C_i the generic creature that is i-th in the queue to cross the highway. We denote with $i = 0$, *i.e.* C_0, the creature that is at the crossing point at time t and must decide if to cross or not to cross. Creature C_{+1} is the creature that has decided to cross before the current creature arrived to the crossing point at time t. Unless we specifically provide this information, we make no assumption about the outcome of the crossing of creature C_{+1}. It may have been successful or unsuccessful. We call C_{-1} the creature next in line after C_0 and C_{-2} the creature next in line after C_{-1}, *etc.*, see Fig. 2.

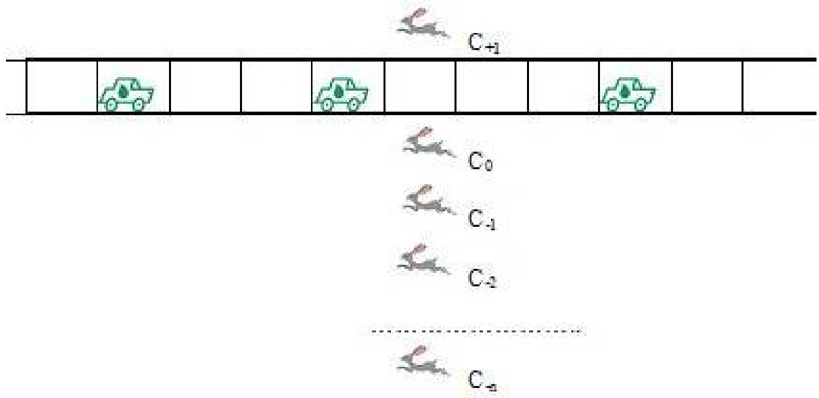

Fig. 2. Simplified diagram showing the highway, several cars, and several sequentially identified agents.

We call time t the time when the creature under consideration must make a decision if to cross or not to cross, we assume that this creature has observed the incoming traffic at times $t-1$, …, $t-n$, where the value of n depends on experiment. This observation allows the creature to evaluate the approximate speed of an incoming vehicle within a set of possible categories, *i.e.* very slow, slow, average, fast, and very fast. We assume that the creature remembers what has happened to other creatures who have attempted earlier crossings in similar circumstances. As long as the creature is capable of this type of observation, it can build a mental table with all possible outcomes for all possible combinations of vehicle distance and vehicle velocity from the cell where the crossing may take place, see Table I. If crossing results in a creature being struck all other creatures will "write" a "1" for the specific distance velocity combination, while if the crossing is successful a "0" will be left for the specific distance velocity combination. In other words, the mental table in the beginning is populated with 0s in the assumption that all possible distance velocity combinations allow crossing. After a certain number of crossings the table may have be modified to look as in Table II, where the combination of high speed and proximity have been found to be hazardous to the creatures who have crossed.

TABLE I

Table showing all possible combinations of vehicle position (*i.e.*, rows) and vehicle velocity (*i.e.*, columns) at the beginning of a new experiment (*i.e.*, no prior knowledge is assumed and no leaning has taken place).

	Very fast	Fast	Average	Slow	Very slow
Very close	0	0	0	0	0
Close	0	0	0	0	0
Mid range	0	0	0	0	0
Far	0	0	0	0	0
Very far	0	0	0	0	0

TABLE II

Hypothetical table content shown after a certain number of crossings.

	Very fast	Fast	Average	Slow	Very slow
Very close	1	1	1	0	0
Close	1	1	0	0	0
Mid range	1	0	0	0	0
Far	0	0	0	0	0
Very far	0	0	0	0	0

A potential problem is due to the approximate nature of the linguistic concepts adopted for defining speed and distance. For instance, it is conceivable that while the combination of "Mid Range" and "Average" is safe most of the times it may result in an accident at some other time. It is thus reasonable to expect that after a certain number of safe crossings for the combination "Fast" (speed) "Mid Range" (distance), it may happen that a creature may be struck for the same crossing. In fact, the actual values within a range may be different and result in a different outcome. Considering the linguistic labels as fuzzy sets, one can apply all the tools available within fuzzy logic and all paradigms of "Computing with Words". They can be found elsewhere; see for instance [9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24,25,26,27,28] and [29]. However, rather than trying to determine when it may be safe to cross, one may also assume that the population of creatures is not very intellectually sophisticated and that it is actually very conservative and prudent. Thus, once a combination of "speed" and "distance range" has been found to be unsafe, all creatures will avoid it.

2.2. Two lanes one directional highway

In [3] we wrote that *a hierarchical layered model should be defined for the generic cognitive agents in a style akin to the hierarchical OSI 7 layer model used in data communication.* We added that *If the interface between layers is kept sufficiently generic, one could achieve an interoperability similar to the one possible with the OSI 7 layer model in data communication.* The first layer of this architecture was meant to be a *Perceptual Layer, akin to the physical layer in the OSI 7 layer model.* If we map the two lanes in such a way that they are perceived as one lane by each creature, there will be no need to modify the creature and the accumulated knowledge will be applicable to the new situation. As a matter of simple example, this can be achieved by logically OR-ing the traffic in one lane with the traffic in the other lane and by transferring this information to the creature.

2.3. Two lanes on a bi-directional highway

The situation when the creature may have to cross a bi-directional highway with traffic flowing in opposite directions is more complex. It may be decomposed into two consecutive crossings of one lane unidirectional highways as in Sec. 2.1. The creature can:

- assess the situation with regard to half of the highway, for instance with traffic flowing eastward;
- cross the highway, if safe, with traffic flowing eastward;
- stop at the boundary between the two halves of the highway;
- assess the situation with regard to the other half of the highway, with traffic flowing westward;
- cross the highway, if safe, with traffic flowing westward;
- repeat the sequence.

Clearly, we cannot just lift a generic creature from one artificial world and place it into another artificial world and expect that it will operate correctly. However, if we design the artificial worlds in such a way that all relevant information can be correctly transferred from the environment to the agent (*i.e.*, the creature), we can have a realistic agent capable of adapting to different environments within a certain catalogue of possible environments. One does not need to think in advance of all possible environments and does not need to design software for all of them. By implementing the software using the object oriented paradigm, it is possible to implement new entities as derived classes, thus inheriting the behaviours and the properties of existing classes and developing only what is new.

3. Conclusion and future work

We believe that our contribution consists in having designed a very simple yet effective abstraction and related model. We are currently designing the software implementation of the proposed agent. We feel that by carefully designing the decision table, *i.e.* Table I, it is possible to provide an agent with a primitive level of "reasoning" and "judging". Moreover, the granularity of the decision table, *i.e.* the number of classes of distance and velocity, may affect the goodness of the model. We will design the agent in such a way that it can operate unchanged for different number of lanes and different direction. We will test this design with several simulated experiments.

REFERENCES

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

[2] B.N. Di Stefano, A.T. Lawniczak, IEEE Proc. of Science and Technology for Humanity (TIC-STH), 2009 IEEE Toronto International Conference, Toronto, Sept. 26–27, 2009, pp. 509–514.

[3] A.T. Lawniczak, B.N. Di Stefano, *Elsevier Procedia Computer Science* **1**, 2221 (2009).

[4] G. Indiveri, E. Chicca, R.J. Douglas, *Cognitive Computing*, Springer Science+Business Media, LLC 2009, 9 pages (published on-line Jan. 14, 2009).

[5] A. McIlroy, *The Globe and Mail*, A3 (Feb. 22, 2011) (available online at `http://m.theglobeandmail.com/news/technology/science/artificial-intelligence-pioneer-aims-to-make-computers-learn-like-brains/article1915447/?service=mobile`).

[6] ISO/IEC 7498-1:1994(E), `http://www.iso.org/iso/catalogue_detail.htm?csnumber=20269`

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

[8] S. Maerivoet, B. De Moor, *Phys. Rep.* **419**, 1 (2005).

[9] L.A. Zadeh, *IEEE Comput. Intell. Mag.* **3**, 11 (2008).

[10] L.A. Zadeh, *Inf. Control* **8**, 338 (1965).

[11] L.A. Zadeh, *Inf. Sci.* **8**, 199 (1975); **8**, 301 (1975); **9**, 43 (1975).

[12] L.A. Zadeh, *Synthese* **30**, 407 (1975).

[13] L.A. Zadeh, *Fuzzy Sets Syst.* **1**, 3 (1978).

[14] G.J. Klir, B. Yuan, *Fuzzy Sets and Fuzzy Logic: Theory and Applications, 1/e*, Prentice Hall PTR, 1995.

[15] H.J. Zimmermann, *Fuzzy Set Theory — And Its Applications*, Third Edition, Kluwer Academic Publishers Group, 1996.

[16] K.M. Passino, S. Yurkovich, *Fuzzy Control*, Addison Wesley Longman, Menlo Park, CA, 1998.

[17] B. Kosko, *Fuzzy Engineering, 1/e*, Prentice Hall, 1996.

[18] B. Kosko, Proc. of the First IEEE Conference on Fuzzy Systems (FUZZ-92), San Diego, March 1992, 1153–62 (1992).

[19] B. Kosko, *Fuzzy Thinking — The New Science of Fuzzy Logic*, Hyperion/Disney Books, 1993.

[20] J.A. Dickerson, B. Kosko, Proc. of World Congress on Neural Networks (WCNN-93), **2**, 9 (1993).

[21] J.A. Dickerson, B. Kosko, *IEEE Trans. Syst. Man Cybern. Part B* **26**, 542 (1996).

[22] B. Kosko, S. Mitaim, *IEEE Trans. Fuzzy Syst.* **9**, 637 (2001).

[23] B. Kosko, *Neural Networks and Fuzzy Systems: A Dynamical Systems Approach to Machine Intelligence*, Englewood Cliffs, NJ: Prentice Hall, 1992.

[24] N. Kasabov, *IEEE Comput. Intell. Mag.* **3**, 23 (2008).

[25] H. Ishibuchi, K. Nozaki, N. Yamamoto, H. Tanaka, *IEEE Trans. Fuzzy Syst.* **3**, 260 (1995).

[26] K.M. Chow, A.B. Rad, *Fuzzy Sets Syst.* **132**, 147 (2002).

[27] S. Mitra, Y. Hayashi, *IEEE Trans. Neural Networks* **11**, 748 (2000).

[28] L.A. Zadeh, *IEEE Trans. Fuzzy Syst.* **4**, 103 (1996).

[29] L.A. Zadeh, *IEEE Trans. Circuits Syst. Fundam. Theory Appl.* **4**, 105 (1999).