

Cognitive Agents for Microscopic Traffic Simulations in Virtual Environments

Sven Seele¹, Rainer Herpers^{1,2,3}, and Christian Bauckhage⁴

¹ Bonn-Rhine-Sieg University of Applied Sciences, 53757 Sankt Augustin, Germany
{sven.seele,rainer.herpers}@brsu.de

² University of New Brunswick, Fredericton, E3B 5A3, Canada

³ York University, Toronto, M3J 1P3, Canada

⁴ University of Bonn, 53115 Bonn, Germany
bauckhage@bit.uni-bonn.de

Abstract. Traffic simulations in current open world video games and driving simulators are still limited with respect to the complexity of the behavior of simulated agents. These limitations are typically due to scarce computational resources, but also to the applied methodologies. We suggest adding cognitive components to traffic agents in order to achieve more realistic behavior, such as opting for risky actions or occasionally breaking traffic rules. To achieve this goal, we start by adding a personality profile to each agent, which is based on the “Five Factor Model” from psychology. We test our enhancement on a specific traffic scenario where simplistic behaviors would lead to a complete standstill of traffic. Our results show that the approach resolves critical situations and keeps traffic flowing.

Keywords: virtual environments, multi-agent systems, cognitive agents, traffic simulation.

1 Introduction

Recently, open world games have become very popular, since they allow players to roam freely through large areas and enable a non-linear playing experience. Obviously, this experience also requires virtual worlds to be populated by believable agents. If these simulated agents use means of transportation such as automobiles, simulations of vehicular traffic need to be realistic. However, state-of-the-art traffic simulations for digital games usually only display relatively simple, rule-based behaviors, so they do not mimic the reality of road traffic. This hampers the player’s feeling of immersion. In serious games applications, such as driving simulators for road safety education (cf. the FIVIS bicycle simulator system [5]), a lack of realism is in fact harmful.

Our goal is therefore to equip virtual agents with skills and reasoning mechanisms that produce real-world behaviors. With respect to traffic simulations, the critical observation is that, in everyday traffic, rules are frequently broken, either on purpose or because of human error. The aim of simulators, such as in

the FIVIS project, is to teach children to identify and resolve potentially hazardous situations resulting from misconduct. If simulated traffic participants are to be useful for road safety training, they need to have a tendency to make errors and to occasionally defy traffic rules. However, in order for the misbehavior to make sense to an observer, it must be grounded in the momentary circumstances. Hence, the agents must be able to perceive and reason about their environment before making decisions. Unfortunately, current traffic simulations typically rely on scripted or randomized behaviors and therefore rarely produce realistic hazardous situations. In this paper, we therefore investigate the effects of adding psychological profiles to agents.

2 Related Work

Traffic simulations have been researched since the 1930s (see [14] for an exhaustive overview of different approaches). The field is mainly concerned with investigating emergent traffic behaviors in order to plan roads or predict jams. While applications like these require general models of driving behavior, they are not concerned with individual habits or actions of agents. In virtual environments, however, where users directly interact with a simulated world, individual behaviors are a major component of the user's experience.

Traffic simulations for virtual environments are often of limited complexity, typically because of scarce computational resources but also because of the current state of artificial intelligence. In fact, some researchers have criticized the limited character of game AI as one of the major shortcomings of current digital experiences and have focused on providing adaptive game AI using methods such as case-based reasoning or reinforcement learning (see e.g. [2][12][13]).

Research on cognitive architectures naturally influences work on modeling cognitive traffic agents. While there are numerous different paradigms and concepts [14][9], ICARUS [10] is an example of a corresponding framework applied to the domain of in-city driving. ICARUS enables an agent to learn how to perform basic driving tasks, for example, aligning itself in a lane or accelerating or decelerating for turns. Similarly, the work in [11] used a cognitive architecture to model erroneous behavior of pilots and drivers. However, it focused on attention mechanisms for assistance systems and did not address decision-making.

Earlier work in [8] presented the concept of a more realistic traffic agent based on psychological personality profiles. Its core is the "Five Factor Model" (FFM) which states that every personality can be defined using only five distinct personality traits (openness, conscientiousness, extraversion, agreeableness, neuroticism) [3][6]. However, the correlation between an agent's personality and its behavior in traffic was not worked out. Nevertheless, the FFM provides an auspicious starting point for the work presented here.

3 Modeling Cognitive Traffic Agents

While the approaches reviewed above apply many interesting concepts, they do not address the problem of modeling *misbehavior*. However, to achieve a realistic




Fig. 1. Cognitive traffic agents for virtual environments could contain the depicted modules from cognitive architecture research. The psychological profile is depicted in a separate layer to indicate that it may influence all cognitive processes.

traffic simulation, agents need to behave and misbehave just as humans do. When they misbehave, they should only do so in situations where it makes sense to an observer, thus ruling out solutions solely based on randomization. We therefore aim to devise agents that are willing to intentionally take certain risks to reach their goals. Choosing risky actions could cause the agent to break traffic rules, which in turn would create potentially dangerous situations for other traffic participants. Our hypothesis is that modeling cognitive processes performed by human drivers and equipping agents with the ability to imitate these processes will result in the desired realistic behavior.

The first important aspect is to model what an artificial traffic participant should be able to perceive of its environment. As human drivers have limited perception, realistic agents need to be equipped with a “perception filter” that controls which information actually reaches the agent. We therefore model the visual field of an agent geometrically within a 3D game engine. Other agents residing within this geometric construct are considered to be perceivable. The concepts of attention selection and attention division presented in [11] can be added in a modular fashion to refine this simple perception model. Figure 1 shows a corresponding example of a cognitive traffic agent. It draws heavily on existing cognitive architectures and features typical elements of this domain, for instance, a short and long-term memory, decision and learning processes, as well as links to the agent’s environment in the form of perception and action modules. Additionally, emotions, mood and feelings can influence the current perception/attention of the agent as well as decision-making.

To ground the agent’s possible cognitive processes in realistic traffic behavior, we add a psychological component to each agent. The “Five Factor Model”, a common method in psychology, is used to classify personalities based on five distinct character traits. Researchers have tried to correlate the model’s factors with driving behavior, but according to Herzberg the results show varying success, with conscientiousness being the only trait consistently connected to driving

behavior [6]. However, in his studies, Herzberg was able to show a correspondence between driving behavior and personality prototypes. The prototypes are based on trait configurations and three classes are especially prominent: resilient, overcontrolled, and undercontrolled drivers.

To identify a subject’s personality, tools like the NEO-Five-Factor Inventory (NEO-FFI) can be used. The NEO-FFI is a questionnaire that assigns scores to the five traits of the FFM [3]. Individuals belonging to the resilient class score low in neuroticism but high for all other traits. *Overcontrollers* score high in neuroticism and conscientiousness, low in extraversion and openness, and average in agreeableness. *Undercontrollers* show high scores in neuroticism and openness, low scores in conscientiousness and agreeableness, and average scores in extraversion. Drivers belonging to the class of *undercontrollers* tend to take more risks and drive more aggressively. On the other hand, *overcontrollers* tend to closely follow traffic rules. *Resilients* are individuals whose behavior lies somewhere in between the other two classes [6].

To test our agent model, we defined a scenario which challenges regular rule-based traffic agents. If agents simultaneously approach an intersection from all sides, the rules of right-hand traffic demand that each agent yields to the one coming from the right. If every traffic participant abided strictly by the rule, a deadlock situation would result, leading to an irresolvable congestive state. To prevent this, drivers in real traffic would communicate – typically through hand signs – who is willing to waive the right of way. Interestingly, one of the participants would thus consciously break a traffic rule but resolve the deadlock.

To determine which traffic agent gives up its right of way in a deadlock situation, we considered the three psychological prototypes. Since *undercontrollers* tend to drive aggressively, they are not likely to give in. This explains the low agreeableness scores of this class. People with high scores in agreeableness usually oblige to other people’s needs, tend to avoid confrontation, and adapt their behavior. High agreeableness scores among *resilients* suggest that drivers from this group would probably allow other drivers to pass. One would suspect overcontrolled drivers to do the same, since they are overly careful. Yet, they also strictly adhere to traffic rules and might thus be more reluctant to waive their rights. This is well reflected in the fact that *overcontrollers* score only average in agreeableness. Since *overcontrollers* and *resilients* are more likely to yield than *undercontrollers*, and since *resilients* are more likely to yield than *overcontrollers*, we use a combination of conscientiousness and agreeableness traits to determine which traffic agent waives the right of way in a deadlock situation.

4 Experimental Setup

To investigate the proposed ideas, a model of an uncontrolled intersection with a priority-to-the-right system in effect, was implemented within the Unity 3D engine as shown in Figure 2. The basic system consists of a closed loop of about 1 kilometer of road around a single intersection in the center, i.e. each lane leads to and from the intersection. The system is populated with a specified number




Fig. 2. For evaluation, a closed road system with an uncontrolled four-way intersection was modeled. Colored vehicles represent psychological prototypes: *undercontroller* (red), *overcontroller* (blue), and *resilient* (yellow). In the depicted situation a resilient driver waives its right of way (indicated by a yield sign) allowing the driver on its left to cross the intersection which resolves the deadlock situation.

of traffic agents which follow the road until they reach the intersection. There, they randomly decide whether to go right, left or straight on.

For the basic driving behavior, the intelligent driver model (IDM) was implemented. The IDM is a simple yet complete model for microscopic traffic simulations [7], which considers maximum acceleration, desired velocity or minimum headway. The parameters and their values were chosen according to suggestions in the literature [14]. At the core of the IDM is the calculation of the agent's current acceleration. In a free-flow driving state the agent simply accelerates, based on the given preferences and vehicle characteristics, until it reaches its desired velocity. In the experiment, the desired velocity was set to 30 km/h which typically is the speed limit in areas with priority-to-the-right intersections. If the agent is following another vehicle, a deceleration term based on the difference between the current headway and the desired headway results in adaptations of velocity.

The first part of the experiment was to fill the road network with a specified number of simple, rule-based agents. The agents were randomly distributed across the network and once the simulation started, the time taken to reach a congestive state was measured. For each specified number of agents (15, 20, 25, 30, and 35) fifteen trials were performed.

For the second part of the experiment the rule-based agents were enhanced with a component representing a psychological personality profile. Agents were divided into *overcontrollers*, *undercontrollers*, and *resilients*. Since Herzberg found an approximately even distribution of the three classes amongst drivers [6], we also distributed our simulated traffic agents evenly across personality prototypes. Initial values for agents were set according to the NEO-FFI scores of the respective psychological traits, scaled to a range between -1 and 1.

In a deadlock situation, the personality traits of the agents determine who gives up its right of way. To implement this mechanism, agents were given the ability to perceive the information necessary to detect deadlock situations and to signal their willingness to waive their priority. To test the robustness of the system, several simulations with 35 vehicles were run for 15, 45 and 120 minutes.

5 Results

Assuming an average vehicle length of 4 *m*, theoretically the maximum capacity of the entire road network corresponds to about 530 vehicles. As expected, deadlock situations occurred for numbers far below this value. The exact moment when a deadlock occurs, depends on the initial vehicle distribution and routing decisions at the intersection. Thus, the fewer number of agents in the system, the higher the chance that a deadlock situation in the system will be delayed. Nevertheless, even at only 3% of maximum capacity (15 vehicles), deadlocks were observed after only 333 seconds on average. At the same time, we also observed two occasions where the simulation continued for 1172 *s* and 957 *s*. However, since the median value (216 *s*) was considerably lower than the average, these cases can be considered as outliers. As illustrated by Figure 3, the time taken to reach a congestive state rapidly decreased as the number of agents in the system increased. An increasing number of vehicles, too, diminished the effect of advantageous routing decisions, since more agents had to wait at the intersection.

Introducing only five additional agents into the simulation decreased the average deadlock-free time by about 3 minutes, while still leaving room for times above the 10-minute mark. However, at 4.7% capacity (25 vehicles), the average and median times fell below one minute, with mean values of 43, 36, and 28 seconds and median values of 46, 36, and 26 seconds for trials with 25, 30, and 35 vehicles respectively. In accordance with the aforementioned observations, the values depicted in Figure 3 also show a decrease in deviations of simulation times as the number of agents increases.

Using our extended model, the simulation was run for three different periods of time. In each of the three trials the traffic never came to a complete standstill. 42 deadlock situations occurred during the 15-minute trial, 116 during the 45-minute trial, and 327 during the 120-minute trial. All situations were resolved allowing the traffic to keep flowing. Figure 4 illustrates a direct comparison of the rule-based agents and our extended model. To achieve comparability between simulations, the random number generator was initialized with identical values,




Fig. 3. The box plots depict the times until the simulation of rule-based agents resulted in a deadlock situation at the four-way intersection. For each number of vehicles 15 simulations were run.




Fig. 4. A comparison of rule-based agents with the extended model using psychological personality profiles. The simulations were run for 10 min. with the first deadlock occurring after 45 sec. (vertical line). The flow (num. of vehicles per min.) was measured at the intersection (a) and the average speed of all vehicles in the simulation (b).

resulting in equal decisions for vehicle distribution and routing. Then the traffic flow at the intersection was measured by counting how many vehicles crossed the intersection in a time interval of one minute. The first deadlock occurred after 45 seconds, but as evident from Figure 4 (a), the flow never reached zero for the enhanced agents. At the same time, the average speed of each vehicle in the simulation was recorded for the duration of the simulation. Figure 4 (b) shows how the average speed continuously decreases after the first occurrence of a deadlock situation for both agent types. However, for the enhanced model the average speed never decreases below 10 km/h .

6 Conclusions and Future Work

In this paper, we discussed the need for more realistic traffic agents in games and simulations. We suggested extending regular, rule-based traffic agents with cognitive components. In particular, in an attempt to equip agents with more realistic behavior, we suggested the addition of psychological profiles based on the “Five Factor Model”, dividing agents into three prototypical groups. We evaluated the impact of this extension in a specific traffic situation that cannot be handled by rule-based agents, and we found that our approach is able to seamlessly resolve deadlock situations at uncontrolled intersections.

At the current state of development, the static personality of an agent determines which agent gives up its right of way. In the worst case, this would always be the same agent, forcing it to wait for an extended time to cross the intersection. Currently, we are working towards modifying static personalities by dynamic elements such as emotions, mood and feelings. In this way, frustration levels could be modeled to increase with prolonged waiting times, which would in turn decrease the likelihood of agents giving up their right of way too often. Other improvements concern the agents’ perception and the number and complexity of their available actions. However, the most challenging task remains

to implement models of internal perception and planning to increase the realism of the agents' overall behavior for traffic simulations in virtual environments.

Acknowledgements. The FIVIS project received additional funding from the Deutsche Gesetzliche Unfallversicherung (grant FP307). The AVeSi project is being funded by the FHprofUnt program of the BMBF (grant 17028X11).

References

1. Anderson, J., Bothell, D., Byrne, M., Douglass, S., Lebiere, C., Quinn, Y.: An integrated theory of the mind. *Psychological Review* 111, 1036–1060 (2004)
2. Bakkes, S., Spronck, P., van den Herik, J.: Rapid and reliable adaption of video game ai. *IEEE Transactions on Computational Intelligence and AI in Games* 2(1), 93–104 (2009)
3. Borkenau, P., Ostendorf, F.: *Neo-Fünf-Faktoren Inventar (NEO-FFI) nach Costa und McCrae*. Hogrefe, Göttingen (1993)
4. Byrne, M.: Cognitive Architecture. In: Sears, A., Jacko, J.A. (eds.) *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*, pp. 97–117. Lawrence Erlbaum, Mahwah (2003)
5. Herpers, R., Scherfgen, D., Kutz, M., Bongartz, J., Hartmann, U., Schulzyk, O., Boronas, S., Saitov, T., Steiner, H., Reinert, D.: Multimedia Sensory Cue Processing in the FIVIS Simulation Environment. In: Ghinea, G., Andres, F., Gulliver, S. (eds.) *Multiple Sensorial Media Advances and Applications: New Developments in MulSeMedia*, pp. 217–233. IGI Global (2011)
6. Herzberg, P.: Beyond “accident-proneness”: Using Five-Factor Model prototypes to predict driving behavior. *J. of Res. in Personality* 6(43), 1096–1100 (2009)
7. Kesting, A., Treiber, M.: Agents for traffic simulation. In: Uhrmacher, A.M., Weyns, D. (eds.) *Multi-Agent Systems: Simulation and Applications*, pp. 325–356. CRC Press, Boca Raton (2009)
8. Kutz, M., Herpers, R.: Urban Traffic Simulation for Games. In: *Proc. of the Int. Conf. of Future of Game Design and Technology*, pp. 181–185. ACM, New York (2008)
9. Laird, J.: Extending the Soar cognitive architecture. In: Wang, P., Goertzel, B., Franklin, S. (eds.) *Proc. of the 2008 Conference on Artificial General Intelligence*, pp. 224–235. IOS Press, Amsterdam (2008)
10. Langley, P., Choi, D.: A Unified Cognitive Architecture for Physical Agents. In: Cohn, A. (ed.) *Proc. of the National Conf. on AI*, pp. 1469–1474. AAAI Press (2006)
11. Lüdtke, A., Weber, L., Osterloh, J.-P., Wortelen, B.: Modeling Pilot and Driver Behavior for Human Error Simulation. In: Duffy, V.G. (ed.) *ICDHM 2009*. LNCS, vol. 5620, pp. 403–412. Springer, Heidelberg (2009)
12. Sharma, M., Holmes, M., Santamaria, J., Irani, A., Isbell, C., Ram, A.: Transfer learning in real-time strategy games using hybrid CBR/RL. In: *Proc. of the 20th Int. Joint Conf. on AI*, pp. 1041–1046. Morgan Kaufmann, San Francisco (2007)
13. Thureau, C., Bauckhage, C., Sagerer, G.: Learning Human-Like Movement Behavior for Computer Games. In: *Proc. of the Int. Conf. on the Sim. of Adaptive Behavior*, pp. 315–323. MIT Press (2004)
14. Treiber, M., Kesting, A.: *Verkehrsdynamik- und Simulation*. Springer, Heidelberg (2009)