SEEKING AND AVOIDING COLLISIONS
A Biologically Plausible Approach

M. A. J. Bourassa and N. Abdellaoui
Defence Research and Development Canada Ottawa, 3701 Carling Ave., Ottawa, Ontario, K1A 0Z4, Canada

Keywords: Agent models, Collision avoidance, Proportional navigation.

Abstract: The success of an agent model that incorporates a hierarchical structure of *needs*, required that the needs trigger human-like *actions* such as collision avoidance. This paper demonstrates a minimalist, "rule-of-thumb" collision avoidance approach that performs well in dynamic, obstacle-cluttered domains. The algorithm relies only on the range, range rate, bearing, and bearing rate of a target perceived by an agent. Computation is minimal and the approach yields a natural behaviour suitable for robotic or computer generated agents in games.

1 INTRODUCTION

An agent architecture based on Maslow's *Hierarchy of Needs* (Maslow, 1943), was introduced by Bourassa et al.(Bourassa et al., 2011) as a possible approach to obtaining more human-like behaviour from computer-generated characters. The approach was inspired by the belief that purely reactive agents that populate many games cannot display human like behaviour because they lack a basic human trait: motivation.

To varying degrees, human behaviour is modulated by motivating factors other than a stated goal or rules. An agent model should therefore incorporate some aspect of motivation.

Maslow proposed a theory of human motivation (Maslow, 1943) that was *human-centered* and founded on *the integrated wholeness of the organism* and *goals* as opposed to *drives*. It was postulated that humans possessed a hierarchy of needs. The most basic needs were physiological: food, drink, rest, etc. As these needs are satisfied, other needs emerge successively: safety, love, esteem, and self-actualisation. Figure 1 illustrates a proposed agent model incorporating Maslow's Hierarchy of Needs. A more detailed illustration of the concept of agent needs is provided in Figure 2. A one-to-one correspondence must be established between the needs in Maslow's Hierarchy and equivalents for a generic agent hierarchy with the nature of the agent being the guiding factor. For example, human hunger and thirst may correspond to fuel and lubricant levels in a robotic agent but to "energy level" in a computer generated character in a video game.

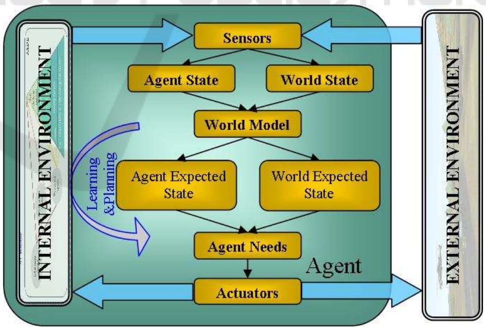

Figure 1: A proposed agent model based on Maslow's Hierarchy of Needs. The agent has sensors and a world model for two environments: internal and external. The agent's actions are decided upon by the fulfillment of its *needs* (see text).

Each level in the hierarchy is comprised of several individual needs. Each need is represented by a *goal* and *utility*. Goals are common to all needs and levels, and map a need to an action. Utilities are unique to each need and provide a model of a need that maps to a level of satisfaction that, in turn, modulates the action mapped to a goal.

In the simplest implementation of the preceding agent model, any action taken is in response to the lowest unsatisfied need. The model provides human-like decisions. The actions themselves, however, must be human-like if a credible semblance of human behaviour is desired. This paper therefore focuses on creating human-like behaviour for one action. This

A. J. Bourassa M. and Abdellaoui N..
SEEKING AND AVOIDING COLLISIONS - A Biologically Plausible Approach.
DOI: 10.5220/0003753602400245
In *Proceedings of the 4th International Conference on Agents and Artificial Intelligence* (ICAART-2012), pages 240-245
ISBN: 978-989-8425-96-6

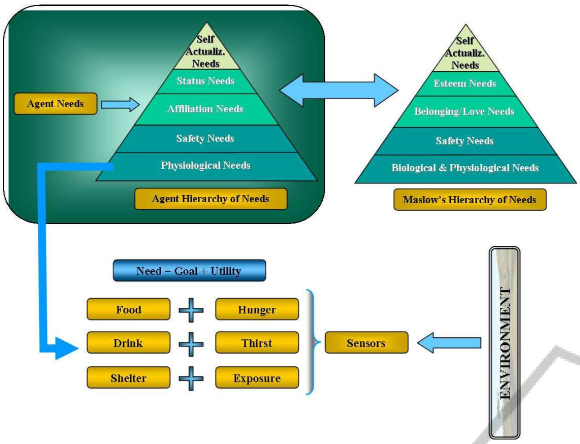

Figure 2: The proposed Agent's Hierarchy of Needs. The lowest unmet needs in the pyramid are the highest motivators. Each need is associated with a *goal* and *utility*.

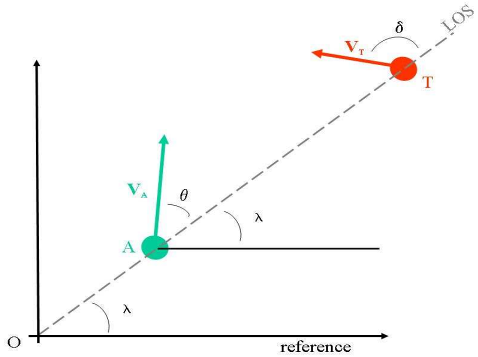

Figure 3: Guided by an observer at *O*, The agent begins at *O* and attempts to intercept *T* by maintaining a course along the LOS (dashed line).

paper addresses one aspect of the lowest level of Maslow's Hierarchy, the Physiological Needs invoking the action of collision avoidance. The implementation of collision avoidance is fundamental to artificial organisms such as non-player characters (NPC's) in computer games and simulations. Collision avoidance is currently implemented either by using extensive geometric algorithms or very simplistic and fixed sets of rules (Millington, 2006). The result is collision avoidance at the cost of predictable, mechanistic behaviour.

In this paper, an algorithm for collision avoidance for mobile autonomous agents is defined based on principles known to be employed in nature. A known biological technique for collision avoidance, *Constant Bearing Decreasing Range* (CBDR) is explored to fulfil two purposes: collision avoidance and its converse, intercept. Taking a theoretical approach centred on the agent instead of the more traditional absolute reference system (Shneydor, 1998), gives insight into rules-of-thumb for both interception and collision avoidance. The rules-of-thumb yield a simply implemented natural behaviour, not necessarily an optimal solution. The goal is for error-tolerant, not error-free, behaviour. The rules do not depend on assumptions of predictable behaviour by the targets or obstacles, making them ideal for dynamic environments.

2 PROPORTIONAL NAVIGATION

The principle of CBDR states that a collision will occur with any object that: remains on a steady bearing relative to one's own direction of motion and has a range that is decreasing. The principle is the foundation of *parallel*, or *proportional*, navigation (PN). PN is used in missile guidance, by predators in the animal world (Ghose et al., 2009), and by humans for driving or catching balls (Shneydor, 1998).

Figure 3 shows a common representation of the intercept problem. Assume a planar engagement where a moving missile (agent or robot), *A*, seeks to intercept a moving target (or goal), *T*. The angle, λ, is the *line-of-sight* (LOS) between the agent and target measured with respect to the reference. θ is known as the *lead angle* of the agent, while δ is the *path angle* of the target.

Simplifying assumptions to facilitate an intercept solution are that the target moves at constant speed in a known direction. Successful interception can only occur when the agent adjusts course to remain on the LOS and overtakes the target. Naturally, as the target moves, λ changes and so the agent must adjust its path angle, $\theta + \lambda$, in order to remain on the LOS.

LOS navigation is applicable in a situation where robot guidance is remotely provided by an observer. In this sense, it is a "three-point" guidance scheme appropriate for, say, radio controlled robots. The agents of interest for this paper react autonomously to their sensor inputs. In this sense, they have a "two-point" guidance scheme.

An autonomous agent does not have third-party guidance so the relative motion between the agent and an object with which it may collide, is the most important measure. One reason why a self-referential view is more valid for an autonomous agent is that an absolute reference may not always be available nor relevant. For example, a third-party observer may not have complete information of both agent and target due to an obstruction. Second, agent sensors operate relative to the agent and so inherently provide such information. The translation of the principles of LOS guidance to a two-point guidance scheme is depicted in Figure 4. An agent "A" (green) and a target

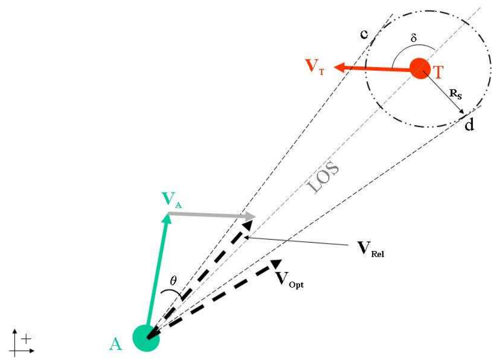

Figure 4: A simple system of typical variables for a collision avoidance or interception in a rotating frame of coordinates. For Proportional Navigation, all variables are assumed to be known hence it is possible to calculate the relative velocity, V_{Rel}, and seek an optimal desired velocity, V_{Opt},that lies outside the *collision cone* ($\triangle_{Acd}$).

"T" (red) are shown in an arbitrary dynamic relationship similar to that of Figure 3. The agent is moving with velocity V_A and the target with velocity V_T. The agent attempts to maintain the LOS at an angle θ to achieve interception of the target. An optimal solution is possible under the assumption of perfect information.(Shneydor, 1998)

Collision avoidance is the opposite of interception where a change of θ is *induced* in order to *avoid* interception. The relative velocity is determined and, if it lies within the "collision cone"[1], a lateral acceleration is provided to redirect the vector outside the cone, V_{Opt}.

The "proportional" in PN is that to effect a change of relative velocity, the agent induces a lateral acceleration proportional to the rate of change of θ. Therefore the classical solution for PN guidance (also known as True PN (Shneydor, 1998)) is:

$$a_{A_L} = -KV_C\dot{\theta} \qquad (1)$$

where K is the *navigation constant*. Explicit values for K can be found when constraints are put on target manoeuvrability and velocity, otherwise one can consider it an arbitrary constant. There are several slight variants on the equation.(Shneydor, 1998; Zarchan, 1994)

3 PROBLEM OUTLINE

PN can be applied to collision avoidance by reversing the guidance laws that apply for intercept, that is, *inducing* $\dot{\theta}$ to generate a miss. It is typically assumed an agent/robot will be in an intercept (sometimes referred to as "navigation") mode while proceeding to a goal and then will switch to a collision avoidance mode in the presence of obstacles meeting some criteria. The principal challenges in using PN have been: dynamic and/or multiple obstacles/targets and high computational demands.

With the exception of Menon's work on missile guidance using fuzzy PN (Menon and Iragavarapu, 1998), approaches to robotic navigation assume: that precise information is available to the agent, that computational demands can be met, and/or that simplifying assumptions of target motion are valid. These are reasonable assumptions for aircraft (manned or otherwise) with sophisticated sensor suites, but not for small or more poorly equipped vehicles. Biological organisms are unlikely to have such perfect knowledge of targets. In this paper the following assumptions are made:

- target range, bearing, range rate (V_C), and the bearing rate ($\dot{\theta}$) are known within an error;
- target aspect and velocity are unknown and no assumptions are made on either; and
- lateral acceleration is not applied explicitly but indirectly through either a change in velocity or bearing.

These assumptions are biologically plausible and sufficient to allow a robust version of proportional navigation.

4 RULES-OF-THUMB

In Figure 5, the geometry of the guidance problem is recast to a purely agent-centric point of view. In such a framework, proportional navigation distills to balancing the velocities of the agent and the target that are perpendicular to the LOS (essentially the intent of *True Proportional Navigation* or TPN (Shneydor, 1998)). Additionally, no assumptions are made, or required, about target behaviour, orientation, or velocities. Since the velocities perpendicular to the LOS of the agent and target must remain equal for interception to occur, the agent must first detect when the velocities are not equal and then apply some means of compensation. The challenge is that V_T, and hence $V_{T\perp}$, is not known.

In the first case, the agent needs to assess whether there is a change in bearing angle, $\dot{\theta} > 0$, or range, $V_C > 0$. This mandates that the agent have at least two consecutive "looks" at the target to obtain rate

[1] The collision cone is a triangle defined by the position of the agent (vertex) and the points of intersection of two lines tangent to a circle representing a "safe" radius, R_S, about the target.

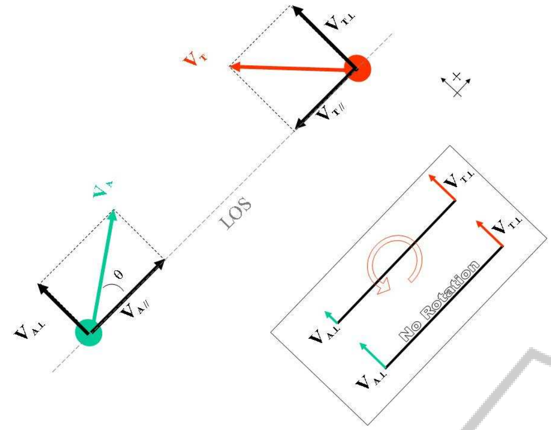

Figure 5: In this simple system, guidance is considered within a agent-centric frame of reference. A line-of-sight (LOS) between an agent (A) and target (T) is defined with respect to an agent's heading. From this point-of-view, a collision will occur if the velocity components perpendicular to the LOS remain equal and the parallel components sum to a positive value.

information. This will provide instantaneous rates of change for both bearing and range.

In the second case, one must identify what options there are for an agent to effect a change in $V_{A\perp}$. The following expresses the rate of change of the perpendicular velocity in terms of agent heading and velocity.

$$\begin{aligned} V_{A\perp} &= V_A \cdot \sin\theta \\ \frac{dV_{A\perp}}{dt} &= \dot{V_A} \cdot \sin\theta + V_A \cdot \cos\theta \cdot \dot{\theta} \end{aligned} \quad (2)$$

Changes to $V_{A\perp}$ can be effected by varying agent velocity or applying a change of heading. The equation (2) highlights the influence of the magnitude of θ. At $\theta \sim 90$ degrees, the $\sin\theta$ term dominates, while at angles nearer to 0 or 180 degrees, the $\cos\theta$ term dominates. This implies that at times, changes in V_A have more effect on $V_{A\perp}$ while at others, changes in heading have a greater effect. When combined with PN principles, this insight yields simple rules-of-thumb that can be used by an agent to produce a very natural behaviour for collision avoidance and interception.

A first rule-of-thumb is the CBDR:

- *A constant LOS bearing and closing velocity indicate a collision is imminent.*

This rule is the one that precipitates one of two courses of action. In a collision avoidance mode, this rule will signal that inducing a change in LOS bearing is required. If the desired action is interception, then this rule signals that no further changes are required to the agent's current state.

Other rules-of-thumb derived from Equation 2 are:

- *If the LOS bearing is changing, and the instantaneous LOS bearing is small ($0° \leq \theta << 45°$), compensate by changing heading.*
- *If the LOS bearing is changing, and the instantaneous LOS bearing is large ($45° << \theta \leq 90°$), compensate by changing velocity.*
- *If the LOS bearing is changing, and the instantaneous LOS bearing is approximately 45°, either a change in heading or change in bearing will suffice. The deciding factor would be whether the agent has the energy or ability to increase V_A.*

A less obvious rule that governs which direction to turn to change heading is:

- *For intercept: adjust heading in the same direction as the bearing rate change. Do the opposite for collision avoidance.*

This means that for interception an agent will steer in direction to minimize $\dot{\theta}$ while for collision avoidance, the intent is the opposite. This appears counter-intuitive as it will mean that collision avoidance will require possibly briefly steering *towards* the target's general direction.

In all cases, Equation 1to determine the magnitude of bearing and velocity changes is used.

5 IMPLEMENTATION AND RESULTS

The agent model (Figure 1) is destined to be used in the OneSAF (Systems, 1998) and other Computer Generated Forces (CGF) environments (Bourassa et al., 2011). A simpler platform was used for algorithm development consisting of Netlogo (Tisue and Wilensky, 2004) and the R Programming Language (R Development Core Team, 2011).

Netlogo is a 2-D, multi-agent, programmable modeling platform used to create agents, render the agent environment, and run simulations. It provides, and manages, agent sensors and agent interactions. Netlogo does not have a full suite of computational libraries and so, in anticipation of future work, the *R Programming Language* was coupled to Netlogo via an extension (Thiele and Grimm, 2010). The rules were then coded in R scripts and called by the agents in the Netlogo environment.

A multi-agent environment was created with 2-30 agents each with a different initial course and speed. The agents were divided into two groups with each group having a goal that ensured that the groups must cross paths enroute to their respective goals as well as avoid fixed obstacles. Each agent was programmed:

- to proceed towards its goal (intercept using equation 1) unless there were other agents or obstacles within sensor range;
- if other agents or obstacles were within sensor range, assess whether the nearest represented a collision danger; and
- apply rules for collision avoidance (the negative of equations 1).

The assessment of collision danger followed the first rule of thumb. The nature of the Netlogo environment meant that bearing rates of zero could not occur. Implementation of the rule was done with various threshold values for bearing rate over a range of $\pm 0.2 - 1.0$.[2] Each agent had a "vision cone" and values of $\pm 10°$ to $\pm 90°$ were tried. These values were chosen as roughly comparable to human vision. Additionally, 10% Gaussian noise was added to all measurements meaning that all agents operated with imperfect information as might be expected in a biological organism. Finally, each agent had a maximum imposed on speed, acceleration, and bearing rate. These substituted for factors such as turning-rate or vehicle dynamics.

Figure 6 illustrates a simple collision scenario involving six agents. Each of the maneuvering agents has altered course to steer behind an oncoming agent. Note that this involved no path planning nor was any perfect information available to any agent about other agents. The agents were further challenged with less linear motion by introducing fixed obstacles. This caused agent movement to be nonlinear and unpredictable. Figure 7 shows agents threading their way past obstacles and other agents while proceeding to goals far above and far to the right of the screen capture. The algorithm called for intercept of the goal in the absence of obstacles thus the agents altered between goal-seeking and collision avoidance. Finally, a more congested view shown in Figure 8 shows more intricate manoeuvering as well as the challenges of a simulated environment. Here fourteen agents pursue distant goals. The centre-most agent has threaded its way effectively past an obstacle and a green agent. The top, left-most agent has collided with a fixed obstacle but this was an artifact of the way the environment is implemented in Netlogo.

[2] This was one of several "tweaks" of the algorithm imposed by the nature of the simulation environment used. Another was that the division of the "world" in to patches introduces artifacts into sensor performance.

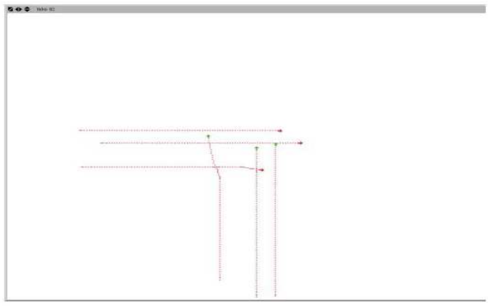

Figure 6: Netlogo screen capture of six agents exercising pure collision avoidance. Each of the maneuvering agents altered course to steer behind an oncoming agent. No path planning was involved. No information of the interfering agents was provided except noisy range and bearing (see text).

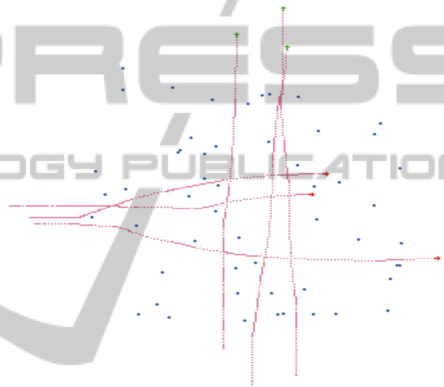

Figure 7: Netlogo screen capture of six agents exercising collision avoidance while proceeding to intercept goals far to the right and above the screen capture. Again no path planning was involved and no perfect information was available. Agent vision was restricted to a 180° cone centred on the agents heading, with a radius of 15 patches.

6 DISCUSSION AND INSIGHTS

The results of the experiments were successful. Despite the simplicity of the algorithms, the behaviour of the agents was natural. For example, agents slowed down at times to allow others to pass, or steered around agents that they overtook. Success with degraded sensor information highlighted that the algorithm is not dependent on high quality or perfect information.

Collisions did occur but this should be viewed in perspective. In the biological world, there are no guarantees of interception nor collision avoidance. A cheetah, for instance, is successful in chasing down prey only 50% of the time (O'Brien et al., 1986) despite an often significant speed advantage; a manoeuvering target is a difficult intercept challenge. Sim-

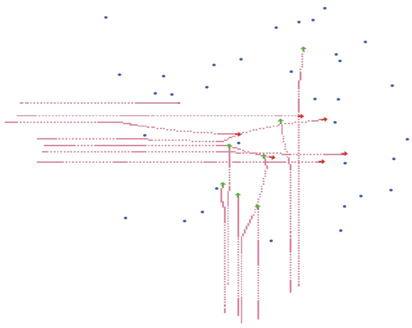

Figure 8: Netlogo screen capture of fourteen agents exercising collision avoidance while proceeding to intercept goals far to the right and above the screen capture. Note the complex manoeuvering of the centre-most red agent.

ilarly, despite rigorous control of airspace and shipping lanes, collisions do occur. The strength of the algorithm is in: its simplicity, its applicability without assumptions or excessive computation, and its robustness to noise.

Target range and bearing rates can be combined with known agent information to derive information about the sensed environment, with simple sensors, and without requiring global knowledge. For example, if T were a stationary object and A were on a fixed course and speed for several time iterations, then T's motion is entirely predictable, that is: it will proceed on a course parallel to A's heading, and it's range rate will behave according to $-V_A \cos\theta$. An even simpler characterization, useful for formation movement, is that any target maintaining the same distance and bearing, $\dot{\theta} = 0$ and $V_c = 0$, is moving at the same speed and heading as the agent.

7 CONCLUSIONS

To satisfy the requirements for an agent model based on motivation, a collision avoidance and interception algorithm was developed using principles known to be used by biological organisms. The algorithm used basic target information obtainable by simple sensors: range, bearing, range rate, and bearing rate. The strength of the approach is that it is simple, robust to noise, computationally undemanding, and biologically plausible. Its implementation is feasible in real time, for real-world platforms with simple sensors. An additional useful insight was the use of range rate and bearing rate to characterize objects detected in the environment.

This work is considered a proof-of-concept and follow-on work in progress includes: using fuzzy logic for rule implementation, implementation in the OneSAF CGF, and an implementation in a mobile robotic platform.

REFERENCES

Bourassa, M., Abdellaoui, N., and Parkinson, G. (2011). Agent-based computer-generated-forces' behaviour improvement. In *ICAART (2)*, pages 273–280.

Ghose, K., Triblehorn, J. D., Bohn, K., Yager, D. D., and Moss, C. F. (2009). Behavioral responses of big brown bats to dives by praying mantises. *Journal of Experimental Biology*, 212:693–703.

Maslow, A. (1943). A theory of human motivation. *Psychological Review*, 50:370–396.

Menon, P. and Iragavarapu, V. (1998). Blended homing guidance law using fuzzy logic. In *AIAA Guidance, Navigation and Control Conference*, pages 1–13, Boston, MA.

Millington, I. (2006). *Artificial Intelligence for Games (The Morgan Kaufmann Series in Interactive 3D Technology)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

O'Brien, S., Wildt, D., and Bush, M. (1986). The cheetah in genetic peril. *Scientific American*, 254:6876.

R Development Core Team (2011). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

Shneydor, N. (1998). *Missile Guidance and Pursuit: Kinematics, Dynamics and Control*. Horwood Series in Engineering Science. Horwood Publishing Chichester.

Systems, L. M. I. (1998). One-saf testbed baseline assessment: Final report. Advanced distributed simulation technology ii, Lockeed Martin Corporation. Commissioned Report for NAWCTSD/STRICOM ADST-II-CDRL-ONESAF-9800101.

Thiele, J. and Grimm, V. (2010). Netlogo meets r: Linking agent-based models with a toolbox for their analysis. *Environmental Modelling and Software*, 25(8):972 – 974.

Tisue, S. and Wilensky, U. (2004). Netlogo: A simple environment for modeling complexity. In *in International Conference on Complex Systems*, pages 16–21.

Zarchan, P. (1994). *Tactical and Strategic Missile Guidance*, volume 157 - Progress in Astronautics and Aeronautics of *AIAA Tactical Misslie Series*. American Institute of Aeronautics and Astronautics, Inc., 2nd edition.