

An integrated perception for autonomous virtual agents: active and predictive perception

By Toni Conde* and Daniel Thalmann



This paper presents an original model with methodologies that integrate in a novel way different types of an autonomous virtual agent's perception in a virtual environment. Our first new approach permits the coherent management of the shared virtual environment for the simulations of an autonomous virtual agent (AVA). Our second approach allows the prediction or the estimation of both the orientation and the attention of an AVA in a virtual environment. By means of a test application with a 'virtual goalkeeper', we demonstrate the speed and the robustness of our technique. Copyright © 2006 John Wiley & Sons, Ltd.

Received: 10 April 2006; Revised: 2 May 2006; Accepted: 10 May 2006

KEY WORDS: perception; virtual sensors; virtual environment; autonomous virtual agents

Introduction

The issue concerning the interface between the decision-taking part and the animating part of an autonomous virtual agent (AVA) is a challenging one. These two parts are represented differently in a virtual environment (VE). For the purpose of visualisation and animation, the VE is pictured by a set of geometric objects; whereas most behavioural models use an abstract and symbolic representation. Perception is an important process that helps reduce the difference between the two parts of the AVA since this enables it to understand its environment. Gilles¹ has further observed that for a simulation of human behaviour to be effective 'it must include the characters' interaction with their environment and to do this it must simulate the character's perception of the environment'.

An AVA uses virtual sensors for *vision*, *audition* and *tactile* to obtain a perception. The AVA has to carry out actions of different degrees of complexity such as evolving in its VE, interacting with it or communicating with other AVAs. To do so, the AVA has to move around or be able to seize objects based on its perception.

Limitations appear when the AVA actions require a dynamic knowledge of the environment, necessitating a perception system. This requirement, together with the intrinsic complexity and cognitive constraints, characterises the boundary between the AVA actions and behaviours. The connection with reflex actions, which require perception but not memory of what has been perceived, takes place at this point. In a common approach, behaviour implements its own perception mechanism, which generates a duplication of calculations, whenever several behaviours are implied.

Contributions

In this paper, we will propose two novel methodologies:

1. The first technique integrates *proprioception* into a *unified* perception concept for an AVA in a *situated* VE. Our motivation was to reach *persistence* and to obtain a *cognitive map* of the perceived VE.
2. The second technique integrates a perception approach by including the faculty of *prediction*, for example, the orientation of the AVA attention.

Related Work

Reynolds² presented a distributed behavioural model for flocks of birds and herds of animals. An interesting result revealed that the 'flocking' behaviour was not only improved by a limited view of the world, but also

*Correspondence to: T. Conde, Virtual Reality Lab—Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland.
E-mail: toni.conde@epfl.ch

Contract/grant sponsor: Swiss National Science Foundation.

depended on it. Tu³ proposed a framework for animation featuring the realistic appearance, movement and behaviour of individual and schools of fish with a minimal input from the animator. Their repertoire of behaviours was reliant upon their perception of the dynamic environment. Individual fish had motivations as well as simple reactive behaviours.

Renault⁴ introduced a synthetic vision system for high-level animation characters. The goal of this system was to let the character move along a corridor whilst avoiding objects and other synthetic characters. Noser⁵ extended previous works by adding memory and learning mechanisms. They considered the navigation feature to be comprised of two parts: a global and a local one. Kuffner⁶ proposed an ethologically inspired approach for real-time obstacle avoidance and navigation. Again, a character rendered the scene from its own viewpoint.

Gilles¹ used a psychological approach to design a visual algorithm. This method did not try to simulate the image on the retina of the character or allow it to perceive features such as colour and shape. Chopra⁷ introduced a framework for generating visual attention behaviour in a simulated human agent based on observations from psychology, human factors and computer vision. Itti⁸ utilised a selective visual attention mechanism, which could rapidly direct the eye towards the objects of interest in the environment.

Peters⁹ proposed a combination of synthetic vision modules with a memory model build on a stage theory from the field of cognitive psychology to let a virtual human attend to its environment. Peters¹⁰ added to their work a system for automatic generation of bottom-up visual attention behaviours in virtual humans. *Bottom-up* attention refers to the way the environment solicits one's attention without regard to task-level goals. Courty¹¹ introduced a model to simulate visual perception based on the detection of salient features, corresponding to the maximum conspicuity in a static image or in a sequence of images. Berthoz¹² advanced the following hypothesis: 'the principle of a huge network of specialised contracted systems, completely decentralised but globally converging, can also be used in an artificial system to integrate different sensorial information and processing algorithms'.

Methodology

Unified Perception

The difficulty of coherence and *unity* of perception cannot be resolved with geometry or dynamics alone. It

also requires the use of central active mechanisms, which can be adjusted for differential delays between sensors.

Several methods have been advanced to implement perception. Our approach proposes a new methodology for *predictive* perception inspired partly by the active perception concepts described by Bordeux¹³ for Synthetic Vision and Database Access. In this model, an AVA maintains a perception puzzle in which each piece corresponds to a specific virtual *sensor*. A *pipeline* is composed of filters capable of extracting relevant information from the data supplied by the related *sensor*. In the next sections, we will describe the main ideas which have been incorporated into our artificial life environment (*ALifeE*) framework.¹⁴ Our approach of a *unified* perception has been integrated in our *ALifeE* framework (see Figure 1). It contains virtual sensors for *vision*, *audition* and *touch* as well as *proprioception*, *active* and *predictive* forms of perception.

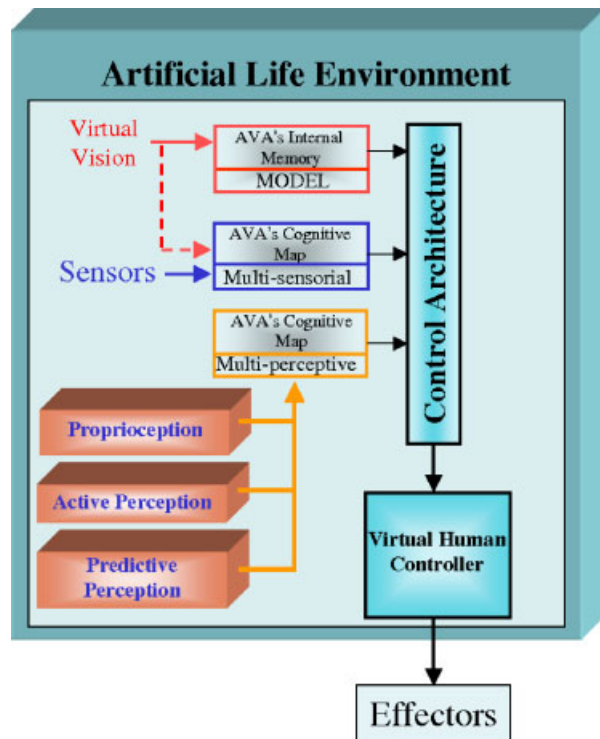


Figure 1. A schematic representation of our ALifeE framework. *Virtual Vision* discovers the VE, constructs the different types of Perception and updates the AVA cognitive map to obtain a multi-perceptive mapping. Then, the control architecture uses the cognitive maps and the memory model to interact with the learning, evolving and control processes of the AVA (virtual human controller).

The information comes from diverse sensorial modalities and is processed by specialised regions with the help of our *ALifeE framework*. This leads to an attempt at a *unified perception*, resulting from the reciprocal connections between these different regions. Moreover, the interaction loops are a particularly effective and quick means of sharing out the processing of information between the various systems, since the entire response time never exceeds that of the slowest system. We have drawn inspiration from the approach described by Ledoux.¹⁵

Proprioception

Proprioception is inspired by the human immune system, which is composed of functional layers combining rapid and archaic mechanisms of innate immunity with slower mechanisms of acquired or adaptive immunity. Our *proprioception* relies on the integration in a single model of three components: *endogenous* and *homeostatic* variables as well as *reinforcement learning* as proposed by Bersini¹⁶ and adapted to our *ALifeEframework*:

1. The notion of *endogenous* variables relates to the internal state of the AVA, which is affected principally by the AVA perceptions and actions. Additional influences allow differentiation between the effects of similar perception inputs on the resulting AVA actions.
2. The notion of *homeostatic* variables describes variables whose temporal dynamics guarantees that they stay within pre-determined boundaries. Since exceeding

these boundaries would result in either a larger discomfort of the AVA or in the death of the human being, actions are taken to prevent these variables from parting from the set-values.

3. To be truly autonomous, the AVA must not only be capable of intelligent actions, but must also be self-sustaining. The delicate problem of simultaneously maintaining the viability of all the variables was resolved with the implementation of a *reinforcement learning* mechanism, namely Q-Learning Sutton.¹⁷ This provided a sequence of actions to keep the AVA viable despite the strong constraints exerted by the VE and the AVAs own endogenous variability.

Active Perception

In the classical approach, each mode of behaviour is implemented with its own perception mechanism. We have improved the method proposed by Bordeaux¹³ by adding the sensorial modalities of *vision*, *audition* and *touch*. The AVA maintains a set of perception *pipelines*, each one corresponding to a specific virtual *sensor*. A *pipeline* is composed of self-coordinating filters or *AgentFilters*, capable of extracting relevant information from the data provided by the matching *sensor* (see example in Figure 2 for a virtual vision sensor).

Perception becomes *active* once the information has been acquired, filtered and simplified by the different virtual *sensors*. It is then integrated in our *ALifeE framework* (see Figure 1). This perception system can be used with the database of any virtual *sensor* suited to a

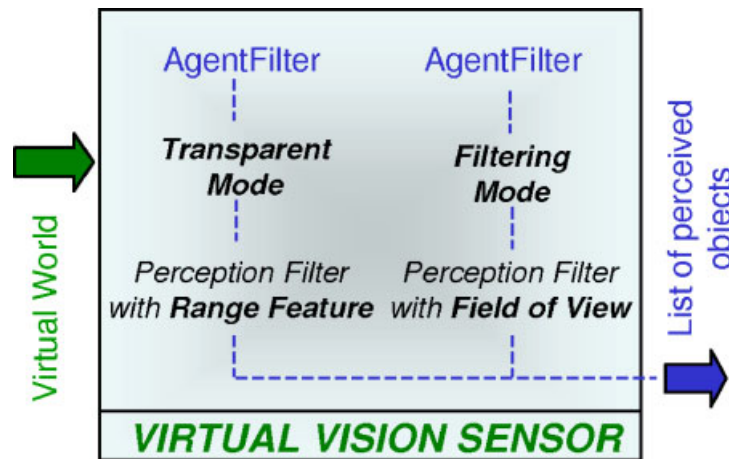


Figure 2. Example of a perception pipeline for a virtual vision sensor.

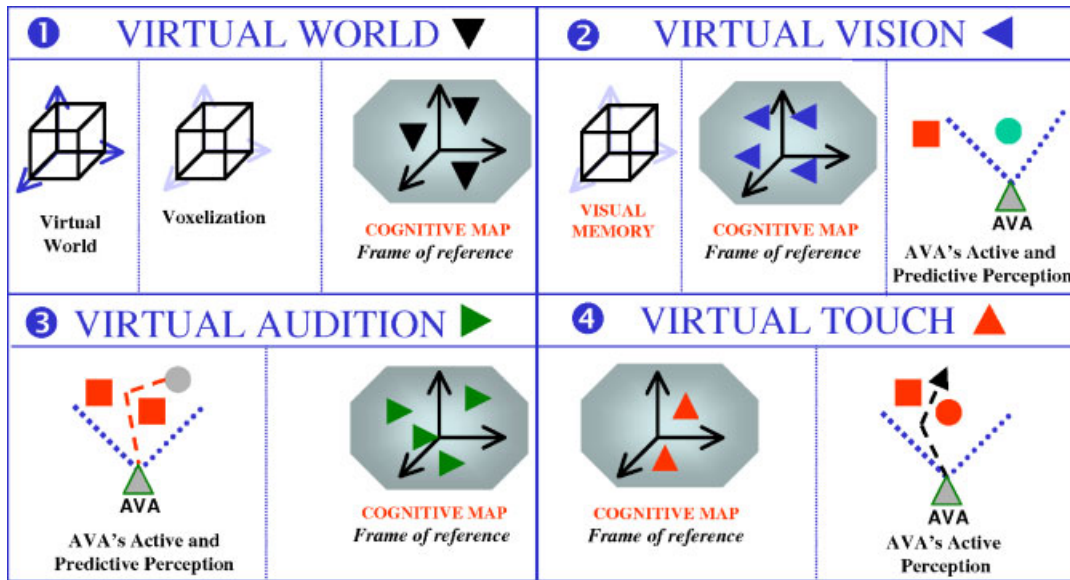


Figure 3. Multi-Sensory Integration. (1) Virtual World—extraction of the information from the VE, voxelization and insertion into the cognitive map; (2) Virtual Vision—extraction of the information from the VE seen by the AVA in coordination with the perception and insertion into the cognitive map and visual memory; (3) Virtual Audition—extraction of the information from the sound environment by the AVA in coordination with the perception and insertion into the cognitive map; (4) Virtual Touch—extraction of the information from the tactile environment seen by the AVA in coordination with the perception and insertion into the cognitive map.

list of perceived sound objects or events (e.g. collisions detection).

Active perception provides the framework necessary for the perception of the AVA behaviour and prepares for a unified perception. For example, the auditory location of an object can be predicted from its visual location through a transformation of eye-centred coordinates to head-centred ones (see Figure 3).

Predictive Perception

An important notion in the AVA perception is the faculty of predicting. Being one of the main activities of the human brain, it naturally plays a fundamental role in active perception. It gives, for instance, the AVA the possibility to direct its look and attention to a sound object.

Our predictive perception model is partly based on the mathematical theory of the observer. Algorithms are used to predict the internal state of a system, generally non-linear, from partial and often external measurements. An observer, a software sensor, is typically composed of a system simulation using an internal, perhaps approximate model, which is guided and

corrected by the measurements taken on the system. In active perception, and under certain circumstances, the observer also permits the selection or the combination of the measurements. Like in the nervous system, it is most useful in improving the estimate of the system state at any given moment. We have drawn inspiration from the functioning of the human eye.

Realisation

Integration of Proprioception

Because of the presence of endogenous variables in the AVA architecture, its behaviour is not simply influenced by its perception Bersini.¹⁶ These variables can be associated with internal mental or emotional states and reactions to external stimuli. They should not be confused with another type of internal variables (form of short memory) which allow the AVA to decide its actions on the basis of not only its current but also past perceptions and actions. These endogenous variables learn to obey homeostatic constraints such as staying within the pre-determined boundaries of a viability zone (see Figure 4).

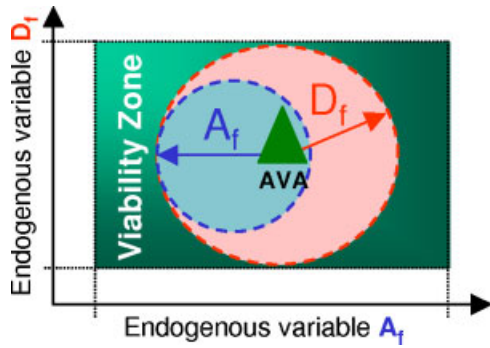


Figure 4. A schematic view of the two endogenous variables A_f and D_f used by predictive perception.

Propioception provides the endogenous parameters necessary to ensure that the predictive perception functions well:

1. Attention Factor (A_f): settles the compromise between an AVA mainly analysing the VE (real-time motion is practically impossible) and an AVA acting blindly (going through objects).
2. Distance Factor (D_f): takes only the obstacles situated inside a defined perimeter into account.
3. Q-Learning: Q-values are updated with the Q-Learning mechanism. The Q-Learning (Equation (1)) learns to select which action to execute for each state (perception state of an AVA P_e and endogenous variables A_f and D_f). At the end of the sequence of actions, the selected action policy maximises a utility function computed over the full sequence with:

$$Q(A_f, D_f, P_e) \leftarrow (1 - \alpha)Q(A_f, D_f, P_e) + \alpha[r + \gamma \text{Max} Q(A'_f D'_f P'_e)] \quad (1)$$

where: α = learning rate and r = reinforcement.

Since the objective for the learning AVA is to maintain its endogenous variables within their viability zone, the type of the reinforcement measure should be chosen accordingly ($\alpha = 0.8$, $r = +1$ if the viability zone is not transgressed).

Integration of Predictive Perception

Our model of predictive perception is based on a set of simultaneous steps:

1. The identification of a specific situation (situated VE) with the use of a database inside the VE.

2. The identification of the object that goes through different physical and contextual mechanisms. The virtual vision sensor provides the identification and approximates the position of the object. The position is represented by a dot, which gives an outline of the borders.
3. The computation of the Priority Factor (P_f) (from 1 to 10)—defines the perception priority to object groups and is fixed in a context (e.g. traffic lights and road signals in a virtual city).
4. The computation of the Attention Factor (A_f) (from 1 to 10)—governs the speed and precision of the analysis in following the object. This factor is provided by the AVA proprioception (see Figure 4). It is variable in the course of time and returns to 1 when the object is identified.
5. The actual identification with the extraction of the useful information.

Our visual part is based on the following hypotheses:

1. The act of vision is undertaken by one eye.
2. The objects move slowly and the AVA_{x,z} moves even more slowly (for the purpose of real-time use).
3. The closer the AVA_{x,z} comes to the object, the greater is the accuracy of perception.

The AVA_{x,z} is located by a position and a vector of absolute direction which is time-dependent. It corresponds to a projection in 3D spherical coordinates (θ , φ and δ) of the AVA_{x,z} vision and to a focalisation plan distance (D_i) with the projection method.

The memorisation of the successive movements depends on both the parameters calculated internally and on the A_f .

The successive steps of modelling for predictive perception involve:

1. Algorithm of visualisation (pseudo-code is illustrated in Figure 5): provides a rough circle (shaded in green) associated to the centre of the object ($Po_{x,z}$) and also displays continuously its outline. Moreover, it implements an approximate field of vision which situates the distance of the AVA_{x,z} from the object (D_i) represented by a semicircle (shaded in orange).
2. Algorithm of space vision (pseudo-code is illustrated in Figure 5): divides the space into concentric circles (coloured in blue) centred on the orientation vector of the observer and displayed continuously. The space is then broken down into areas. Finally, the field of vision is divided into semicircles with a periodicity of radius that decreases as the distance from the AVA

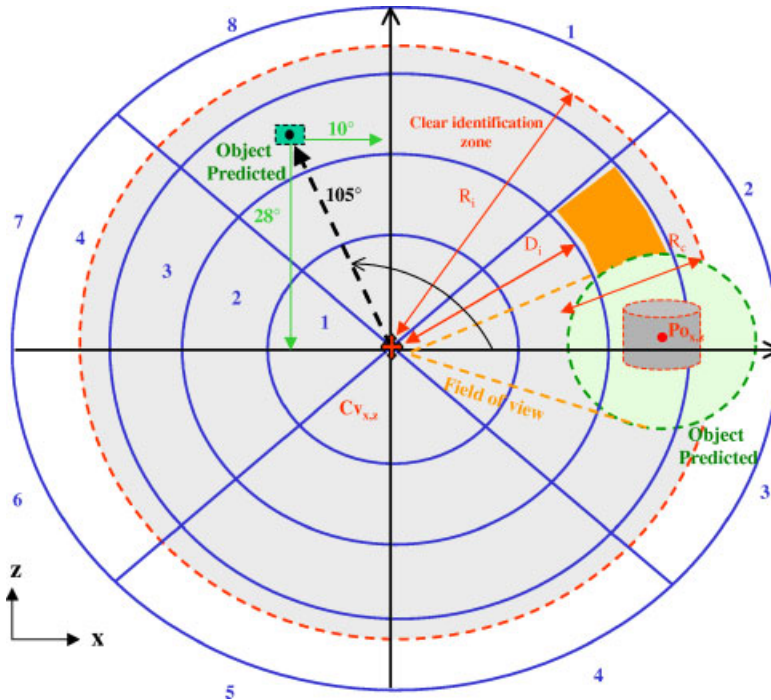


Figure 5. A schematic representation of predictive perception—vision Space of an AVA. R_i , radius of the clear identification zone; D_i , focalisation distance between the centre of vision ($C_{v_{x,z}}$) and the position of the object that should be predicted. R_c , recovery factor between R_i and D_i . $P_{o_{x,z}}$: position of the object that should be predicted.

increases. The operation permits a spatial quantification of the AVA and the A_f is used to increase its level. This allows the model to come closer to human functioning in which perception is approximated with regard to positioning. The model, moreover, is simplified.

3. *Position of object ($P_{o_{x,z}}$):* is associated with the three values necessary for the identification of the portions of the visual space containing the object: a semicircle (shaded in orange), an area in the focalisation plan distance and a field of view with the method of projection.
4. *Static approach—predicted object does not move:* a circle centred on the centre of vision ($C_{v_{x,z}}$) with a variable radius (R_i) defines the zone where the object should be. The radius of the circle is linked to the A_f . To avoid unnecessary movements, the circle associated to the centre of the object and the one associated to the centre of vision must overlap to a certain degree (R_c). This again depends on the A_f . The same concept is applied to the notion of the field of vision. If these conditions are not fulfilled, a movement with a speed and precision determined by the A_f should be induced. The direction of the movement is obtained with the

following formula: projection of the position of the area on axes x and z and normalisation to the defined field of view. As for the focalisation distance, the model is trivial. The object is maintained in the centre of the field of vision until the A_f drops to 1, indicating that the predicted object has been identified.

5. *Dynamic approach—predicted object moves:* if the object is in movement, the successive positions are iterated according to a time constant. The notion of speed is unnecessary as the information perceived by the eye is approximate and corresponds rather to the perception of a movement given by an orientation vector displayed continuously in a virtual simulation.

The number of dots considered and the degree of interpolation are equal. This means that their memorisation starts from 1 and increases, depending on the complexity and the speed of the movement, by a factor that can be calculated. The interpolation allows the measurement of the new position and the deduction of the AVA movement.

The speed of the movement depends on the A_f . The coordinates of the dots are used to induce a movement as in the static approach. The AVA movement produces

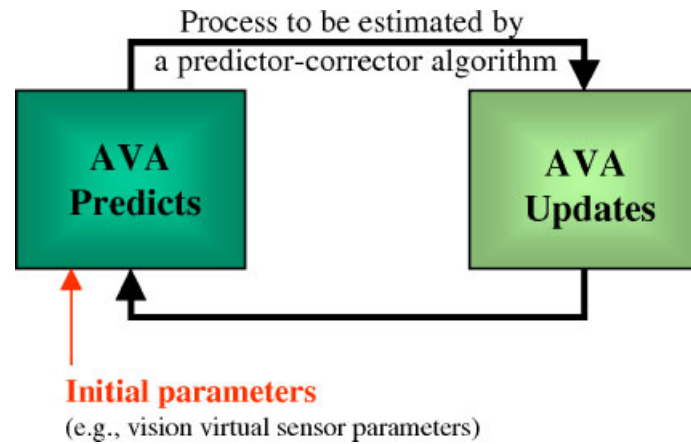


Figure 6. The ongoing discrete Kalman filter cycle. The time update projects the current state estimate ahead in time. The measurement update adjusts the projected estimate by an actual measurement at that time. For specific equations, description of the parameters and the Kalman algorithm, see Reference [18].

a change in the position of the object. The difference between the predicted position and the real one is introduced into our interpolation model. The movement is deduced by an interpolation (see Figure 6) that utilises a Kalman filter.¹⁸

The use of the Kalman filter addresses the following question: given our knowledge of the system behaviour and our measurements, what is the best estimate of the object position? We know how the system behaves (Equation (2)) and we have measurements of the position of the object. So, the Kalman filter permits the evaluation of the state $x \in R^n$ of a quantified time-controlled process that is governed by a linear stochastic difference equation:

$$x_k = Ax_{k-1} + B u_k + w_{k-1} \text{ with a measurement } z \in R^m \text{ that is } z_k = H x_k + v_k \quad (2)$$

The random variables w_k and v_k represent the process and the measurement noise respectively. Matrix **A** relates the state at the previous time step $k-1$ to the state at the current time step k , in the absence of a driving function or a process noise. Matrix **B** relates the optional control input $u \in R^l$ to the state x . Matrix **H** relates the state to the measurement z_k .

Experimental Results

Using different scenarios, we could verify that our *ALifeE framework* associated sensorial and perceptual modalities in a coherent way. Indeed, perception is an

interpretation; its coherence depends on endogenous factors and on the actions. However, perception is above all *predictive*, given that memory helps to anticipate the consequences of a future action.

Therefore, it is not only the number of sensors that is important, but also the questions that the *curious* human brain asks the world based on the hypotheses it elaborates and the tasks it proposes to achieve. The senses constitute a source for hypotheses and function to verify them. Directing one's look is one of the primary functions required by the development of the *curious* brain that simulates action. Berthoz¹² said 'to go in the direction I am looking' and not 'to look where I see'. We simulate the path mentally and we compare the movement that our feet will make with the predicted one. This method has greatly inspired us in the modelling of our *predictive* perception approach.

The Simplified Model

Initially, we adopted a simplified model for the interpolation of the various real and predicted positions of the object. The dot product of the different movements (vectors) of the object that should be predicted was utilised (see Figure 7 and Equation (3)).

$$k = \frac{\sum_{i=1}^n (\vec{a}_i \cdot \vec{a}_{n+1})}{\sum_{i=1}^n d_i} \text{ with } k \in [0.0, \dots, 1.0] \quad (3)$$

where: k = coefficient of the tendency of the movement (trajectory of the object), $\vec{a}_{i...n}$ = positions i of the object and d_i = distance between the positions i .




Figure 7. Simplified model for the interpolation of the movement of the object that should be predicted.

After several simulations in real time, we defined empirically that if the result of the dot product range was between:

1. [0.8...1.0], the trajectories of the object movement were almost *linear*. We used the last vector calculated for the prediction of the object position.
2. [0.4...0.8], the trajectories of the object movement were sometimes *linear* or *jerked*. We therefore used the average of the vectors calculated for the prediction of the object position.
3. [0.2...0.4], the trajectories of the object movement were *strongly jerked*. The prediction was more com-

plex to estimate. We used therefore the Kalman filter for the prediction of the object position.

The Complex Model

The criteria of validation of our method were to predict the position of an object compatible with our real-time constraints, which are specific to the behavioural animation of AVAs.

Figure 8 demonstrates how the Kalman filter handling strongly jerked trajectories of the object movement was able to interpolate the predicted position with the parameters: φ (horizontal angle), θ (azimuth angle) and D_i (focalisation distance), and the real position of the spherical coordinates of the AVA eyes: α (the offset to φ), β (the offset to θ) and δ (the offset to D_i).

After several tests, we considered that a value of 150 Points-Time was sufficient for an AVA to predict an object with jerked trajectories. On Figure 8, the first 150 Points-Time show the evolution of the AVA movements before it can detect the graphic object to predict. Then, during the following 280 Points-Time the AVA carries on its way towards the next graphic object to predict.

The model of interpolation for predicting the object movement was stabilised and integrated into the *ALifeE framework*. Figure 9(a-d) shows successive snapshots of an AVA predicting the movement of a pyramid in a virtual world composed of different graphical objects such as cubes, cylinders and pyramids.

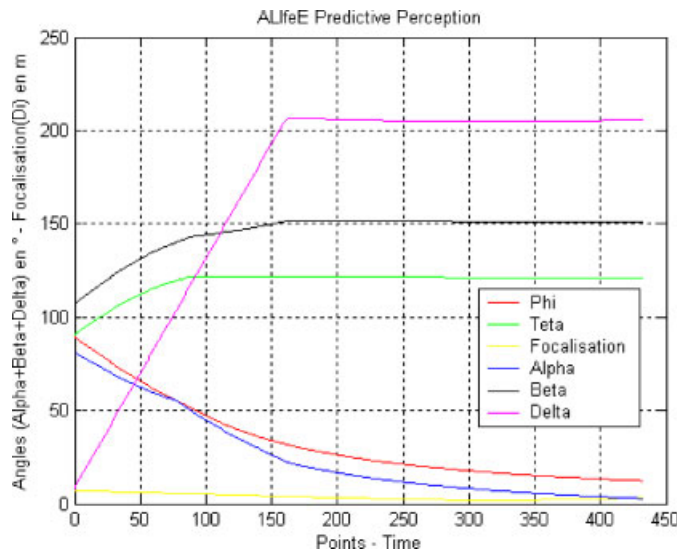


Figure 8. The Kalman filter combined with our predictive perception model was able to correctly estimate the position of the object (measured during 430 Points-Time of the object movement).

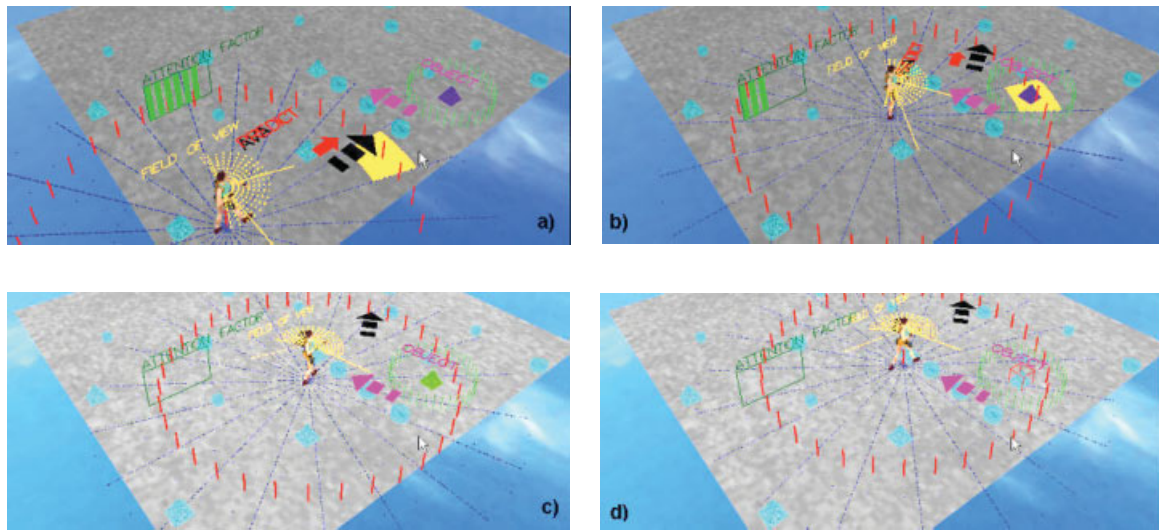


Figure 9. (a–d): Snapshots of ALifeE—AVA predicting the movement of a pyramid. The blue pyramid in (a) and (b) turns green in (c) once it has been predicted. In (d) the AVA carries on its way towards the next graphic object to predict.




Figure 10. (a–d): Snapshots of a 'virtual goalkeeper' with integration of active and predictive perception.

Figure 9(a–d) illustrates the following elements:

- An AVA moving in a VE. The black discontinuous arrow indicates the AVA movement of walk.
- A graphic object represented by a blue pyramid moving in a VE. The purple discontinuous arrow indicates the dynamic movement of the object.
- The AVA field of vision, whose direction is indicated by a red arrow.
- The attention factor (A_f) represented by a green bar-graph.

Test Application With a ‘Virtual Goalkeeper’

To show the feasibility of using active and predictive perception, we have implemented a test application with a ‘virtual goalkeeper’. The AVA (virtual goalkeeper) stands in front of the goal. The objective is to see its different levels of reaction depending on the perceived situation.

The test consists in preventing the ball from entering the goal by placing the hands of the goalkeeper on its trajectory. In order to know if the ball ‘touches’ one of the goalkeeper’s gloves, it is necessary to carry out, at every

moment, a collision detection between the ball and the gloves.

There are two categories of collision detection:

- The first one detects the collision between a moving object and the static geometries belonging to the environment, and
- The second one detects the collision between two moving objects.

It is the latter which we used in our test application, since the ball and the goalkeeper’s gloves are constantly moving on a horizontal plan.

The parameters of our ‘virtual goalkeeper’ are the following:

- The legs of the goalkeeper can move from left to right, and
- The gloves of the goalkeeper, as described above, remain on a horizontal plan.

All animations were rendered in real time using OpenGL on a 3.0GHz PC with an nVIDIA GeForce FX Go5350 video card and are represented in Figure 10(a–d).

On Figure 10(a–d), the ‘virtual goalkeeper’ is represented by a black discontinuous arrow and its

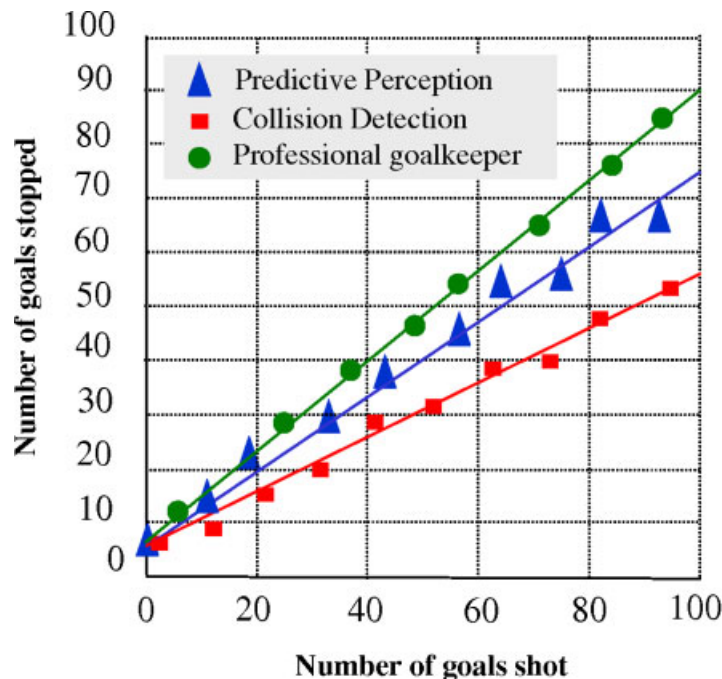


Figure 11. Plot of performances obtained with our predictive perception methodology, with the collision detection approach and with a professional goalkeeper. All animations were rendered in real time using OpenGL on a 3.0 GHz PC with an nVIDIA GeForce FX Go5350 video card.

field of vision is in yellow. Thanks to its active and predictive perception, it is able to perceive the ball whose direction is indicated by a purple discontinuous arrow. The attention factor (A_p) is represented by a green bar-graph.

Comparison With Different Models

We compared in Figure 11 the performance of our complex model with that of the classical approach with collision detection and with the success rate of catching penalties by a professional goalkeeper provided by the Union of European Football Association (UEFA).

Future Work and Conclusions

There are some weaknesses in our approach. Since the *cognitive map* gives only an approximate mapping, entirely correct results cannot be guaranteed. Also the *cognitive map* inputs must be processed with care. Our methodology in case of parallel movement of both the AVA and the object can introduce instabilities. The problem of observational instabilities can be optimised by introducing asymmetries in the boundaries between the different zones. This will be the objective of our next research.

Our approach is part of a more complex model. The goal is to realise a virtual life environment for an AVA including different interfaces and sensorial modalities coupled with various evolving learning methodologies.

ACKNOWLEDGEMENTS

This research has been partially funded by the Swiss National Science Foundation.

References

1. Gilles M. Practical behavioural animation based on vision and attention. *University of Cambridge Computer Laboratory, Technical Report TR522*, 2001.
2. Reynolds CW. Flocks, herds and schools: a distributed behavioural model. *Computer and Graphics* 1987; 4 (21): 5–34.
3. Tu X, Terzopoulos D. Artificial fishes: physics, locomotion, perception, behaviour. In *Proceedings of SIGGRAPH* 1994; pp. 43–50.
4. Renault O, Magnenat-Thalmann N, Thalmann D. A vision-based approach to behavioural animation. *Journal of Visualization and Computer Animation* 1990; 1: 18–21.

5. Noser H, Renault O, Thalmann D, Magnenat-Thalmann N. Navigation for Digital Actors based on Synthetic Vision, Memory and Learning. *Computers and Graphics* 1995; 1: 7–19.
6. Kuffner J-J, Latombe J-C. Fast synthetic vision, memory, and learning models for virtual humans. In *Proceedings of Computer Animation, IEEE*, 1999; pp. 118–127.
7. Chopra S, Badler N. Where to look? Automating attending behaviours of virtual human characters. In *Proceedings of Autonomous Agents and Multi-Agent Systems*, 2001; 4(1/2): pp. 9–23.
8. Itti L. Visual attention. In *The Handbook of Brain Theory and Neural Networks*, (2nd edn.), Arbib MA (ed.). MIT Press: Cambridge, MA, 2005; 1196–1201.
9. Peters C, O’Sullivan C. Synthetic vision and memory for autonomous virtual humans. *Computer Graphics Forum* 2002; 21 (4): 743–753.
10. Peters C, O’Sullivan C. Bottom-up visual attention for virtual human animation. In *Proceedings of Computer Animation and Social Agents (CASA)*, 2003; pp. 111–117.
11. Courty N, Marchand E. Visual perception based on salient features. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2003; 1024–1029.
12. Berthoz A. *The Brain’s sense of movement*. Editions Odile Jacob, 1997.
13. Bordeux C, Boulic R, Thalmann D. An efficient perception pipeline for autonomous agents. In *Proceedings of Eurographics*, 1999; pp. 23–30.
14. Conde T, Thalmann D. An Artificial Life Environment for Autonomous Virtual Agents with multi-sensorial and multi-perceptive features. *Computer Animation and Virtual Worlds* 2004; 15 (3–4), 311–318.
15. Ledoux J. *The emotional brain*. Editions Suster, 1996.
16. Bersini H. Reinforcement Learning for Homeostatic Endogenous Variables. In *Proceedings of 3rd Int. Conference On Simulated & Adaptive Behaviour*, 1994; 325–333, MIT Press.
17. Sutton RS, Barto G. *Reinforcement Learning: An Introduction*, MIT Press: Cambridge, MA, 1998.
18. Kalman RE. A New Approach to Linear Filtering and Prediction Problems. In *Transaction of the ASME—Journal of Basic Engineering* 1960; 35–45.

Authors’ biographies:



Toni Conde is a senior researcher of the VRLab at EPFL. His research interests include virtual humans in virtual environment, perception and learning. Various executive positions held with Swisscom Group Ltd. including: manager of Consulting and Design, head of department,

principal expert and director of Consulting Services; consultant for various multinational companies, banks and international organisations. IEEE senior member, Computer Society. He has an M.S. degree in computer science and Ph.D. from Swiss Federal Institute of Technology in Lausanne (EPFL).



Daniel Thalmann is professor and director of The Virtual Reality Lab (VRlab) at EPFL, Switzerland. He

is a pioneer in research on virtual humans. His current research interests include real-time virtual humans in virtual reality, networked virtual environments, artificial life, and multimedia. He is coeditor-in-chief of the *Journal of Computer Animation and Virtual Worlds* and member of the editorial board of the visual computer and four other journals. Daniel Thalmann was member of numerous program committees, co-chair, and program co-chair of several conferences including IEEE VR 2000. He has also organised five courses at SIGGRAPH on human animation and crowd simulation. Daniel Thalmann has published numerous papers in graphics, animation, and virtual reality. He is coeditor of 30 books including the recent *Handbook of Virtual Humans*, published by John Wiley and Sons and coauthor of several books. He received his Ph.D. in computer science in 1977 from the University of Geneva and an Honorary Doctorate (Honoris Causa) from University Paul-Sabatier in Toulouse, France, in 2003.