Why to Create Personalities for Synthetic Actors

Paolo Petta, Robert Trappl
Austrian Research Institute for Artificial Intelligence*,
Vienna, and
Department of Medical Cybernetics and Artificial Intelligence,
University of Vienna
{paolo, robert}@ai.univie.ac.at

1 Introduction

The last decade has seen not only the impressive developments from the first brief appearance of a computer-generated character in a feature film[1], over the first digital performer in a movie[2], all the way to the first fully computer-generated theatrical motion picture[3]; synthetic actors have also conquered demanding real-time domains from the (co-)moderation of live television shows to the evaluation of complex real-world scenarios modeled in detail in advanced simulators.

Over this timespan the rendering of the external appearance has progressed from the basic geometrical and physical modeling and animation over the inclusion of dependencies on the simulated environment to event- and object-based behaviours. As synthetic actors were placed in simulated worlds of growing complexity, an obvious requirement that came about was to make them perform *in* these worlds: for convincing realistic behaviour they ought to perform as if they were gathering information as it became available to them and they should interact with objects — e.g., tracking, avoiding, manipulating,...them — in plausible ways. It soon became evident that the first approach of setting out from an omniscient position where all the data of the simulated world was made available to programs driving the actors was very inefficient in a number

*The Austrian Research Institute for Artificial Intelligence is supported by the Austrian Federal Ministry of Science, Transport, and the Arts.

[1]Young Sherlock Holmes, directed by Barry Levinson, visual effects by Industrial Light & Magic.

[2]Casper, directed by Brad Silberling, visual effects by Industrial Light & Magic.

[3]Toy Story, directed by John Lasseter, visual effects by Pixar.

of respects. For instance, those pieces of information that could actually be accessed given the actor's current state in the world had to be laboriously identified and progressively singled out. Similarly, all sequences of actions making up basic behaviours of an agent had to be explicitly prescribed: take as an example an object "catching the eye" of the actor which would then turn around to inspect it more closely. Instead, taking the step forward and equipping actors with virtual sensors, turning their bodies into virtual effectors, and using these devices to guide the execution of behaviours proved to be a superior alternative, with the additional benefit of introducing a well defined interface between the actor and its environment, allowing to deploy actors in different worlds[4].

However, all of these external aspects still cover only part of what makes up an actor: as entities that "act a part or take part in an affair", synthetic actors also have to be endowed with internal capabilities, both affective and cognitive—i.e., in concordance with the following definition by Osherson[10]

> Cognitive Science is the study of human intelligence in all of its forms, from perception and action to language and reasoning. The exercise of intelligence is called cognition. Under the rubric of cognition fall such diverse human activities as recognizing a friend's voice over the telephone, reading a novel, jumping from stone to stone in a creek, explaining an idea to classmates, remembering the way home from work, and choosing a profession. Cognitive processes are essential to each of these activities; indeed, they are essential to everything we do. [p.ix]

From the point of view of the different applications, the actual requirements calling for these internal faculties can vary widely: feature film animators might just want to be able to instruct their puppets using high level commands (e.g., *"walk leisurely across the room"*) or be assisted in the more delicate and ephemeral task of ensuring consistent patterns of behaviour of a virtual actor over a longer time span; in other settings, achieving e.g. various degrees of agent autonomy can play an essential role in providing effective assistance to users or in achieving credible and interesting performance in virtual environments. Especially in this latter context, the inevitable limitations in terms of incompleteness of information about the world that can be perceived by the actor via its virtual sensors as well as restrictions in terms of available resources[5] become of particular relevance. Different approaches have been taken to tackle this problem. Some favour runtime performance and acceptance by a broad public even though adhering to scientifically disproved preconceptions of folk psychology — as in the Oz project directed by Joseph Bates at Carnegie Mellon

[4] A related effort which gained a certain publicity was the work carried out by an undergraduate student at the MIT Media Lab who interfaced their ALIVE system to the game world of DOOM. See [12, 3] for a more detailed discussion of these topics.

[5] Computational resources and timing constraints, especially for characters "thrown" into worlds where they have to act in real time.

University. Others place the emphasis on the soundness of the implementation of a cognitive theory resting on plausible first principles, as illustrated in the chapters by Dave Moffat and Aaron Sloman in this volume. Also, it must not be overlooked that as the exterior architecture of agents and their environments become more complex and structured, control of these worlds necessarily also becomes of a more indirect, higher-level nature, and consequently has to rely on the availability of interior mechanisms of corresponding functionality.

All of these circumstances thus give rise to concerns regarding the characteristics of the high-level behaviour of the synthetic actors, both in terms of reliability and predictability—ensuring that the performance remains within anticipated or prescribable limits—and interestingness—the manifestation of individual distinguishing marks due to which the actors become recognizable individuals: in other words, the design of tailored personalities. The papers included in the present collection span the range from theoretical to applied and application-specific aspects, laying a foundation for the creation of multifarious and more believable synthetic actors. From its very beginning the event at which these works were presented was characterized by an informal atmosphere which engendered animated exchanges of opinions. The results of these discussions are reflected in the contents of the revised versions of the workshop papers. In the following we will briefly summarize some key points presented in the single contributions, deferring a principled discussion of the subject domain to the final chapter of this volume.

2 Contents

2.1 Virtual Humans

An essential point of emphasis of the research supervised by Nadia Magnenat-Thalmann and Daniel Thalmann at the MIRALab in Geneva and the EPFL in Lausanne is the pursuit of the highest qualitative standards (i.e., photorealism) of the appearance of their creations: achievement of flexible real-time behaviour is subordinated to this primary goal. Consequently, special consideration is given e.g. to the simulation of different surfaces, such as clothes made of different fabrics and materials, or human hair and skin [7]. Similarly, single courses of actions—such as the grasping of objects or the facial expressions—are modeled in accurate detail [12]. These latter procedures also rely on the last focus of research to be mentioned here, namely the synthetic sensing of the environment. In particular this includes the virtual vision, tactile, and audition, which contribute decisively to a "natural" behaviour of the virtual actors. Most of this research is strongly application-oriented, with telecooperation—following the maxim that the realistic representation of Avatars[6] has a crucial impact

[6] This commonly adopted designation for figures representing human users in cyberspace was first introduced in the Habitat System [9].

of the acceptance of the system—and the entertainment industries (e.g., the co-moderation of prime time shows) as important application areas.

The research efforts presented by Bonnie Webber and Norman Badler [1] also direct their attention on realism: the work centered on *Jack*® focuses on ergonomic (e.g., of complex controls systems such as a helicopter cockpit) and logistic (e.g., in military domains) evaluations of scenarios. Jack comprises a whole family of variations on a basic theme, the biomechanical simulation of a human body in real time. The overall design of Jack follows the three-level-architecture which was popularized in robotic research: on the lowest level a *sense-control-act* loop provides the interface to the environment and ensures robust "reactive" behaviour, making Jack capable of e.g., "stumbling" over obstacles[7]. At the highest level, the behaviour of all instances of Jack in one running simulation are controlled by *parallel transition networks*, in which all possible courses of action for a given scene are defined in an abstract manner (roughly comparable to Roger Schank's *scripts*): to solve occurring partial problems, specialized problem solving methods can be activated as required (e.g., heuristics and planning procedures to chose a hiding place in the context of a game of hide-and-seek).

At New York University's Media Research Lab (NYU MRL) an alternative approach is taken for a similar subject [4]. In the framework of a procedural animation technique, single movements of human figures are defined along with parameter settings for all possible transitions between these actions. As evaluation criterion for this latter step the assessment of the "naturalness" of the resulting impression as given by human experts was drawn up. This exact modeling is subsequently softened by the superposition of stochastic noise functions, so that each repetition of a single movement is performed with slight variations; but also when at a standstill, the joints of the figures thereby feature a slight "unrest", which contributes to the realistic appearance.

Given that in this way the single movements can be recalled in an efficient manner, it is possible to control the whole animation at a higher level in real time. This functionality is provided by behavioural scripts, which define the "vocabulary" of each object (actors as well as any other entities comprised in the scene). Instead of equipping each actor e.g. with comprehensive knowledge about all possible actions, the respective relevant information is thereby stored in a distributed fashion (a tea cup "knows", that it can be sipped from, a vehicle knows that it can be mounted and driven, etc.). Among the opportunities arising from this approach, it is also easily possible to achieve the effect of a stepwise acquisition of skills in mastering recurrent tasks by instantiating and

[7] As an interesting side note, the use of reinforcement learning algorithms was found to be of indispensable help in the adjustment of the numerous parameters of this complex numerical system.

updating variables stored with the given actor. The considered application scenarios for this technology include the interactive design of virtual choreographies and virtual social environments akin to graphical Internet MUD (Multi-User Dimensions) and their more recent commercial variants.

2.2 Virtual Ethology

Bruce Blumberg and Tinsley Galyean [3] present further developments of the Hamsterdam architecture which is being applied in the context of the ALIVE (Artificial Life Interactive Video Environment) project at the MIT Media Lab. ALIVE implements the metaphor of a "magic mirror" which allows users to immerse in a virtual environment without requiring them to wear any special equipment: this mirror is realized as a wide-area screen on which the image of the user as acquired by video cameras is displayed after being composited with a computer generated environment. User can manipulate object of the artificial world by means of gestures and similarly interact with its synthetic actors, virtual animals. The Hamsterdam architecture underlying these creatures builds upon Pattie Maes' "action selection" algorithm and particularly also takes into account findings from ethology, thereby achieving interesting, believable, and timely behaviour. Hamsterdam obeys a strict separation of drawn appearance (*geometry layer*), the repertoire of possible movements *motor skills*), and a *behaviour system* which is responsible for the possible behaviours. The behaviour system consists of a loose hierarchy of single behaviours, which are each responsible for the pursuit of a specific goal. Within groups of mutually inhibiting behaviours the *avalanche effect* describe by Marvin Minsky in Mentopolis is used to select a particular member and to achieve an adequate degree of persistence at the same time. The strength of the inhibitory signals is influenced by local variables particular to the single behaviours ("goals and motivations") as well as by input provided by a releasing mechanism. This releasing mechanism acts as a filter that identifies relevant objects and events out of the stream of data supplied by the system of virtual sensors. Actors implemented using the Hamsterdam architecture can be influenced at different levels of control, ranging from vague motivations (by setting single state variables of behaviours) to explicit instructions (imparted to the motor skill layer). More complex directions can be imparted via a combined activation of behaviour groups and the manipulation of the releasing mechanism and virtual sensors.

2.3 Virtual Drama

David Blair and Tom Meyer [2] take on the the topic under investigation from the perspective of a "guided experiencing" of a narrative. The developments which set out from David Blair's feature film "Wax or the discovery of television among the bees" first led to the working up of the material as a hypermedia system containing a number of automated narrating systems. This system was eventually published on the World-Wide Web as "WAXweb" which was also one of the first applications drawing extensively on the capabilities offered by the first release of VRML (Virtual Reality Markup Language). WAXweb was successively integrated with a interactive environment, the HyperHotel text-based virtual environment based on the MOO technology developed at Xerox PARC. In the subsequent projects, "Jews in Space" and "Grammatron" additional steps are being taken—among other by means of the inclusion of active VRML elements—towards the realization of an implicit way of narrating that can be experienced individually: on the occasion the development of semi-intelligent tools for the design and rearrangement of the deconstructed narrative space forms an important part of the projects as well as the realization of implemented operationalisations of existing theories of drama.

The Virtual Theater project at the Knowledge Systems Lab of Stanford University aims at providing a multimedia environment in which users can interact with intelligent agents that function as synthetic actors. By virtue of this additional level of indirection, the agents are not required to *have* specific personality, they only have to be "in character", i.e., display behaviours that stand in agreement with the personalities they are expected to take on in given contexts [5]. Consequently, their work is much more strongly influenced by drama theory, and more precisely by theories of improvisation than by psychological theories. Personalities are modeled only superficially using a small of set of traits that is much more restricted than the one used in the work presented by Brian Loyall, described below. The focus on artistic models of character is reflected in the importance of an effective visualization—both CMU's Edge of Intention and the procedural animation techniques developed at NYU's MRL have been used—and the reliance on a corpus of performing heuristic rules (e.g., how to effectively convey status transactions that are assumed to govern human relationships). This contribution thereby highlights the fact that at least in this particular setting the goal of achieving the impression of a believable performance delivered by synthetic actors embodying distinct personalities can be met with the inclusion of a minimum of "cognitive machinery".

Brian Loyall [6] reports on recent extensions of the well-known "Edge of Intention" scenario developed under the supervision of Joseph Bates at CMU. The possibilities to interact with the Woggles populating this virtual world had long been limited to "body language", i.e. movement and changing of size of the user's Woggle-avatar, which were then interpreted by the other Woggles as friendly or threatening gesture, invitation or readiness to participate in play, etc. Now this repertoire is complemented with the possibility of entering texts or fragments of speech. Woggles in the vicinity of the avatar may react immediately, without necessarily waiting for the completion of the typed input. This engenders the possibility of misinterpretation of the partial information, leading to reactions which have to be corrected later on. For example, an "irritated"

Woggle being addressed may take this action to be a hostile act; an interpretation which may be revised when the content of the message is seen to be a friendly one. In this way, this new means of interacting offers the opportunity to make the emotive inner life of the Woggles—a slightly extended version of the personality theory by Ortony, Clore, and Collins—even better visible to the human user of the system.

2.4 Virtual Cognition

Finally, the contributions by Dave Moffat [8] and Aaron Sloman [11] round off the broad investigation of the topic by providing a psychological-cognitive perspective. Both authors share the view that the traditional methods of psychology fail to provide an adequate framework supporting actual implementations of working "deep" cognitive models.

After reviewing a some of the more important lines of research in the area of personality, Moffat takes the important step of explicitly detaching the concept of personality from the anthropomorphic context assumed in traditional personality theories. He then illustrates how his model roots in Nico Frijda's theory of emotion: this particular theory was chosen as it was seen to be nearer to implementability than most other theories in the field, and furthermore to provide the most comprehensive coverage of the problem domain. Moffat's model, W*ill*, is already a successor to a first partial implementation of Frijda's theory of emotion. A prominent design goal was to keep the model as simple as possible, using off-the-shelf AI technology: "Rather than make a big, sophisticated architecture that models everything, the idea is to include as *few* design features as possible because they correspond to theoretical assumptions, and like many others, I prefer theories ...to be simple, economical, and parsimonious".

Aaron Sloman, on the other hand, takes on a much more radical stance, virtually dismissing all of the hitherto conducted psychological research which, among other things, he accuses to be sorely lacking of such basic essentials as a well-defined terminology. In stark contrast to Moffat's attitude cited above, the complex holistic architecture designed by Aaron Sloman ("Personality belongs to a whole agent") lays claim to be a comprehensive model of cognition, at which the realization of exemplary instances proves to be rather difficult because of the assumed very flexible components (e.g. the open set of representation formalisms and associated virtual machines operating on them, or the complex mechanisms of meta-control) and the *necessarily* rich content that has to be covered. The interpretation of emotion as an emergent phenomenon of *breakdown* following from the limitation of the available resources is just another proposition that will doubtlessly spark many further discussions, such as the ones that took place during the workshop and which are partly reproduced in this volume as appendix to Aaron Sloman's contribution.

3 Concluding Remarks

This brief overview of the topics covered in the contributions comprised in this collection cannot be expected but to convey a first impression of the plethora of facets under which the topic of creating personalities for synthetic actors is being researched. At the same time, within this wide range of diverse approaches pursued to tackle the posed challenges there are evident strong interrelations of which we have tried to point out just a few. This circumstance corroborates the importance of furthering the exchange of ideas between researchers from different backgrounds in this rapidly growing area.

References

[1] Badler N., Reich B.D., Webber B.L. (1997) Towards Personalities for Animated Agents with Reactive and Planning Behaviors. In: Trappl R., Petta P. (eds.) *Creating Personalities for Synthetic Actors* (In this volume)

[2] Blair D., Meyer T. (1997) Tools for an Interactive Virtual Cinema. In: Trappl R., Petta P. (eds.) *Creating Personalities for Synthetic Actors* (In this volume)

[3] Blumberg B. (1997) Multi-level Control for Animated Autonomous Agents: Do the Right Thing... Oh, Not That... In: Trappl R., Petta P. (eds.) *Creating Personalities for Synthetic Actors* (In this volume)

[4] Goldberg A. (1997) IMPROV: A System for Real-Time Animation of Behavior-based Interactive Synthetic Actors. In: Trappl R., Petta P. (eds.) *Creating Personalities for Synthetic Actors* (In this volume)

[5] Hayes-Roth B., Gent R.van, Huber D. (1997) Acting in Character. In: Trappl R., Petta P. (eds.) *Creating Personalities for Synthetic Actors* (In this volume)

[6] Loyall B. (1997) Some Requirements and Approaches for Natural Language in a Believable Agent. In: Trappl R., Petta P. (eds.) *Creating Personalities for Synthetic Actors* (In this volume)

[7] Magnenat-Thalmann N., Volino P. (1997) Dressing Virtual Humans. In: Trappl R., Petta P. (eds.) *Creating Personalities for Synthetic Actors* (In this volume)

[8] Moffat D. (1997) Personality Parameters and Programs. In: Trappl R., Petta P. (eds.) *Creating Personalities for Synthetic Actors* (In this volume)

[9] Morningstar C., Farmer F.R. (1990) The Lessons of Lucasfilm's Habitat. In: Benedikt M. (ed.) *Cyberspace: First Steps*, MIT Press, Cambridge, MA

[10] Osherson D.N., Lasnik H. (eds.) (1990) *An Invitation to Cognitive Science*, MIT Press, Cambridge, MA

[11] Sloman A. (1997) What Sort of Control System is Able to Have a Personality? In: Trappl R., Petta P. (eds.) *Creating Personalities for Synthetic Actors* (In this volume)

[12] Thalmann D., Noser H., Huang Z. (1997) Autonomous Virtual Actors based on Virtual Sensors. In: Trappl R., Petta P. (eds.) *Creating Personalities for Synthetic Actors* (In this volume)