Real-Time Behavioral Animation of Humanoid Non-Player Characters with a Computational Ecosystem

Rui Filipe Antunes and Frederic Fol Leymarie

Goldsmiths, University of London, United Kingdom
{rui.antunes,ffl}@gold.ac.uk

Abstract. A novel approach to a decentralized autonomous model of agency for general purpose Non-Player Characters (NPCs) is presented: Computational Ecosystems as a model of AI. We describe the technology used to animate a population of gregarious humanoid characters in the virtual world *Where is Lourenco Marques?* an ethnographic artistic work characterized as a virtual world inhabited by a population of NPCs interacting autonomously among themselves as well as with an audience of outsiders (human observers). First, we present the background and motivations for the project. Then, we describe the technical details about the algorithm that was developed to generate the movements and behaviors of a population of NPC 'storytellers'. Finally, we layout some of the critical aspects of this particular implementation and contextualize the work with regards to a wider usage in virtual worlds.

Keywords: Multi-agent systems, Simulation, modelling and visualization, Animation, Computational Ecosystem, Virtual Worlds.

1 Introduction

The Ecosystem as an Allegory. To animate a population of Non-Playing Characters (NPCs) we have used a Computational Ecosystem (CE) to play the role of an Artificial Intelligence (AI). This is based on a community of autonomous agents which are organized as a simulation of a food-chain while trading units of energy. Each of the individuals emulates a rudimentary life cycle of generic carbon-based life forms. Mendelian genetics informs the evolution of the community, and genetic characteristics such as the speed or size of an agent are inherited by children from their parents when couples engender in a process that evokes sexual reproduction. Energy is required for the activities these individuals perform in a virtual city, including: moving, running, or simply breathing. The dynamics of energy transfer occur in *predatory* acts when the population competes for energy and space. In particular, when the energy level of an individual is too low (*i.e.* below a pre-specified threshold), this is regarded as death at which point the individual is taken away from the population.

The idea of using a CE to animate a virtual population of humanoids came about to symbolize the social situation experienced in a colonial city. In the tradition of literary allegory, the behavior of the community carries a second, or connotative level of

R. Aylett et al. (Eds.): IVA 2013, LNAI 8108, pp. 382–395, 2013.

narrative. In the virtual city/ecosystem social groups become analogous to *species*. The performance in the habitat dictates the behavior of these storytellers in the virtual world. Nevertheless, instead of the conventional primal events of fierce animals that attack and devour each other, what is shown to the audience/users are animations of humanoids interacting with each other and gesturing in apparent conversation. Each action in the ecosystem is rendered visible in the virtual world as the animation of a corresponding gesture or movement. For instance, if two hypothetical individuals interact, and one "attacks" the other, the expressiveness of the movements of the arms is greater in the winning individual than the movements that its victim shows.

The Computational Ecosystem as an AI Engine. The use of the technology of the CEs as generative engines appears in contexts as diverse as audio-visual installations [1], music [2] or even for the choreography of characters in virtual worlds [3]. The innovation we introduce is the use of the CE as an AI to coordinate the social movements and behaviors of a community of humanoid NPCs, taking advantage of the complexity potential inherent in such a dynamic ecosystem.

By design, the behaviors observed in the community of individuals in a virtual city are translations in the form of visual representations of the inner-workings of the ecosystem. The spatial layout of the virtual city provides a support for the visualization of the underlying multi-agent state space, with a direct connection between the current state of the ecosystem and the configuration of the characters on the layout. Each NPC is an individual belonging to the population and we refer to each one as a 'character'. Each birth is represented by the new character appearing next to its parents, while its 'dematerialization' in thin air represents its 'death'. The interactions during the lifetime of the characters are translated into a series of movements and gesticulations while being constrained to the surface of the world. The CE establishes correspondences between states, movements and actions performed by each of the characters. For instance, the action of an hypothetical individual feeding in the virtual ecosystem might correspond to a certain gesticulation being performed by the character in the virtual world, while its escape after a fight with another creature will correspond to a different gesture being expressed by the character. A set of *nine base animations* were defined for this work. This small set was deemed sufficient for a prototype to allow to explore and demonstrate the potential of CEs in animating virtual populations of humanoid characters. This includes animations for walking (1) and of gestures (8) for the arms to be used during interactions. The implemented mechanism permitting the generation of each character behaviours is described later in § 2.

1.1 Background

The animation of populations in digital reconstructions of historical sites such as Petra, in Jordan, the former Pennsylvania train station in New York city [4], the ancient city of Pompeii [5], the Babylonian Uruk [6], or for theme parks [7], are just a few examples of a flourishing area of research that looks at modeling virtual spaces inhabited by communities of autonomous humanoid characters.

A few standards have been established with regards to modelling groups of NPCs. Three approaches prevail, when individuals: (a) are represented as particles subject to

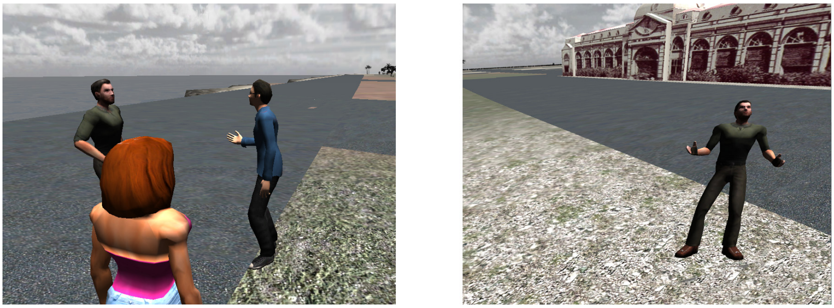

(a) (b)

Fig. 1. (a) Storytellers in autonomous interaction. (b) Storyteller interacting with the audience.

physical forces [8]; (b) are represented as states of cells in cellular automata [9]; or (c) follow a rule-based scheme [10]. One of the merits of these approaches is in modeling in a realistic way at a macro-level the spatial-flow features of a crowd. More recently the individual's behavior within a multitude, at a micro-level, has received the attention of researchers. Some works attempt to recreate the spontaneity of behaviors visible in small groups of a few dozens of individuals. Shao and Terzopoulos define pedestrian behaviors in both (i) the former NYC Penn station and (ii) the theater of Petra's Great Temple, producing NPCs exhibiting heterogeneous and apparently spontaneous behaviors such as standing in queues or watching shop windows [4]. Pelechano *et al.* model a group of humanoids interacting socially at a cocktail party [11]. A good survey of the field of crowd animation up to 2008 can be found in the books by Pelechano *et al.* [12] and by Thalmann and Raupp Musse [13].

We describe later the generative system developed to produce the visual effects of a community of gregarious individuals in small interacting crowds (groups of 2 to 10 individuals, Figure 1). This work was motivated by an artistic piece which required, by design specifications, that in order to address *European colonialism* the animations of a humanoid population inhabiting a virtual city had to be driven by *predatory* behaviors as a metaphor/allegory for the social conditions lived in a colonial city. We are looking, in particular, at systems where agents are organized in the hierarchical structure of a *food-chain* while trading units of energy and biomass as a way of promoting community dynamics. A CE proves to be able to provide for complex environment simulations rich in the heterogeneity and spontaneity which we are targeting.

1.2 The Artwork *Where is Lourenço Marques? (WisLM)*

This paper describes the technical implementation of the CE that animates a virtual population in the artwork *Where is Lourenço Marques?* characterized as a representation of the former city of Lourenço Marques, the capital of the province of Mozambique, during the period of Portuguese colonial domination. The city later became known as Maputo after the independence in 1975. During the process of decolonization from 1974 to 1976, many of its citizens were forced to abandon the city for social and political reasons. This artwork presents the mediation of the memories of some of those who experienced the last period of its colonial times. The memorial takes expression in a 3D virtual world; an illustrative video can be found on YouTube (search: "Where is Lourenco Marques" by "xtnz").

From Interviews to Character 'storytellers'. The initial steps of the project entailed a process of interviews with a community who left the city in those dramatic days, and now mostly resides in Portugal. An artistic and subjective reconstruction of the city was built in a 3D virtual environment (using Unity3D) based on their accounts and their shared material forms of memorabilia. The focus of this paper is the community of humanoid NPCs which roams autonomously the city and interact with each others as well as with the human audience.

These characters become 'storytellers' (Figure 1) when the user selects them (*e.g.* by pointing or clicking on any of them). A selected NPC interrupts its current activity and looks at the camera. Then, while gesticulating expressively, the character provides an audio testimony by streaming one of the oral accounts recorded during the initial process of interviews. Thus, this population of animated NPCs can assist in the task of bringing the past experiences of some of the "expatriated" citizens. Each of the individuals in this population is the bearer of an excerpt from an interview, functioning as the carrier and mediator of real-life human stories. The audience is thus implicitly invited to seek-out the storytellers through the city and listen to their stories.

2 Technical Description

In order to elaborate our CE, we have selected some techniques with a proven record of successful results in animating populations of multiple individuals. These are: (i) a *hormonal system* as proposed by Daniel Jones in his framework for bio-inspired swarms [14], (ii) a *metabolic system* to define and restrict the diet of our characters based on the model of Saruwatari *et al.* [15], and (iii) a *classifier system* adapted from John Holland's model [16] which drives the action–selection mechanism of our characters. We now describe our implementation of each of these main techniques.

2.1 The Hormonal System

Jones' framework for sound-based performance using swarm dynamics, introduces a biomimetic design process that augments the classical rule-set for flocking, in part through the implementation of a hormone-like system which allows temporal modifications of the individual behavior of particles within a swarm [14]. We have adapted this architectural design to expand the classical energy paradigm used in more traditional CEs which typically operates as a thermostat-like regulator [17]. Adding an extra layer of hormonal regulators increases both the life-likeness of our model and its associated complexity. This layer is defined via five variables:

1. Testosterone: increases with age and crowdness, decreases upon giving birth, and also causes an increase in the likelihood of reproduction.
2. Adrenaline: increases with overcrowding, decreases as a result of internal regulation overtime, and also causes a greater rate and variance of movement.
3. Serotonin: increases with 'day' cycles, decreases during 'night' cycles and as result of hunger, and also causes a greater social attraction towards other characters.

4. Melatonin: increases during 'night' cycles, decreases during 'day' cycles and also decreases the rate of movement.
5. Leptin: increases upon eating, decreases steadily at all other times, also causes downward regulation of serotonin when depleted, and finally causes greater attraction to food.

The Blueprint Descriptor. The hormone-like system described above is initially configured by the genetic descriptor which contains a blueprint for the attributes of the character. This is a string with 15 binary digits, where different sections of the string code for a set of six specific features:

1. Age (**gAge**) — defines the rate at which the agent ages.
2. Introspection (**gIntr**) — establishes the level of gregariousness.
3. Hormone cycles (**gCycl**) — the strength or speed of the hormone cycle (*e.g.* how much will it increase per day).
4. Hormone uptakes (**gUptk**) — indicates the intake during an hormone cycle (*e.g.* how much will the hormone increase when a character meets another one).
5. Body chemistry (**gChem**) — defines the chemical components present in the body.
6. Metabolism (**gMetab**) — determines what chemicals can be 'digested'.

In the *initial* population, these features are determined randomly. Afterwards, each subsequent generation inherits this information from their parents. This follows a Mendelian-like process of diploid reproduction, where crossover operators are used on Gtypes to promote variation. In the process some noise is added to the inherited information to mimic mutations (with an arbitrary small probability $Pm = 0.05$ of effective mutation on each bit).

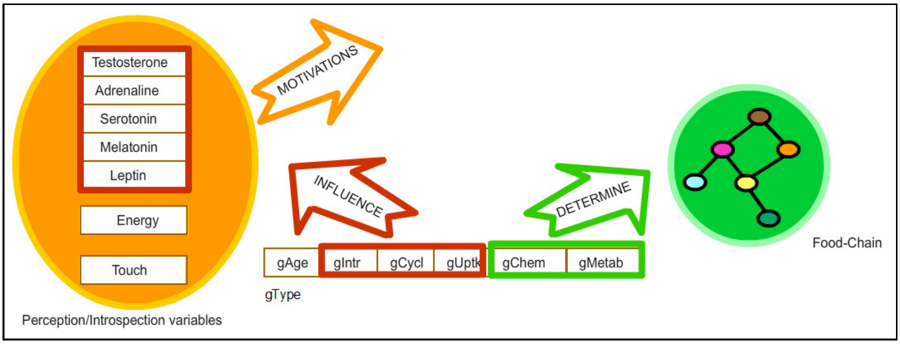

Fig. 2. Graph depicting the GType influence on the behavior of the character. The variables (gIntr, gCycl, gUptk) directly influence the hormonal system, which motivates the actions performed in the world; the variables (gChem, gMetab) set the body-composition and dietary specifications, which determines the environmental context for the actions.

2.2 The Metabolic System

Saruwatari *et al.* [15] provide a model we found useful to determine the dietary specifications. Their framework uses two strings, the first of which defines the body constitution of the character, while the second is used to describe its metabolism. Potential preys are those whose constitution-string matches the predator's metabolic-string. Saruwatari *et al.* have shown this simple mechanism potentially leads to the emergence of complex multi-trophic food-chains with variable depths, which in turn gives us the necessary differentiation and stratification required in our model.

The string of 3 binary digits present in the Gtype section **gChem** defines the first part of the dietary specification representing the character's body 'chemical' composition. These digits code for chemicals A, B and C. Each character is equipped with chemical repositories that are direct translations from the composition-Gtype **gChem**. Take for instance a character with **gChem** of "010". The only chemical present in the repository will be B. On the contrary, another character with Gtype "101" will have both chemicals A and C active. When an hypothetical character (X) preys another character (Y), X will fill its own repositories by extraction from the chemical repositories of Y. In this transfer process 90% of the value is wasted. Each repository only carries a maximum capacity directly related to the character's body size which is given by a direct translation of the binary value of **gChem**. These chemical attributes play a fundamental role in the ecosystem, since they determine part of the dietary specification in the community and thus the interactions between individuals.

The second part of the dietary specification is provided by **gMetab**, the component that defines the character's metabolism, *i.e.* what chemicals can be 'digested'. An example is an hypothetical character with **gMetab** 010, which will be able to prey individuals whose **gChem** codes the B component: 010, 110, 011, 111. Consequently the combination **gChem–gMetab** structures the essential interactions in the predator-prey relationships. This predator-prey mechanism of matching the metabolic-string with the composition-string provides an interaction space of size 8 x 8, which was wide enough for the current work (in terms of observed behaviors).

The Metabolic Rules. The metabolic system simulates a simplified food-energy conversion. Besides the described hormonal system and chemical repositories, another main structuring variable contributes to the behavior of the characters: energy. This is generated from the chemical repository of the character. To emulate a conversion from mass to energy, we have defined an arbitrary chemical reaction which requires three chemical units to produce one unit of energy (*e.g.* $1\,B + 2\,C \Longrightarrow 1\,energy\,unit$). Energy can be spent when breathing or in activities performed in the world such as moving, running away, attacking, playing, mating or eating. Below a pre-set level more energy needs to be produced from the chemical–repositories of the character. Below a pre-set threshold the character starts to feel tired, activating an internal sensor (to search for food). If the energy level reaches the value 0 the character dies and is removed.

2.3 Behavior of the Characters via a Classifier

The characters' behavior is defined by three main operational stages: perception, decision, and action. Each character is provided with: (i) internal sensors monitoring energy and hormonal levels, and (ii) sensors for contact which are triggered by the proximity of other characters (Figures 2 and 3). As a function of these inputs, the character will choose the action to take using a classifier system inspired by the description provided by John Holland [16], a model that allows autonomous agents exhibiting self-organization capabilities and temporal adaptation. Holland's model was used in the well-known Echo system [18] and also inspired artworks such as Eden [17].

The Classifier System. During the process of perception the system inspects the level of energy and the state of the hormonal variables, as well as if the body of a character is entering in contact with any other characters. When any of these variables is activated, such as when: (i) the energy or leptin levels are below pre-selected thresholds, (ii) the testosterone, adrenalin or serotonin levels are above some pre-fixed thresholds, or (iii) the character is touching another body, then an *action–message* is generated. This message takes the form of a string of length 6, and is composed from the grammar set $\{0, 1, \#\}$, identifying which sensor is active — binary values indicate the active and inactive states while # functions as a wildcard.

The active messages list. This is a list with messages queuing to be processed. If a message is new on the list, it will be inserted with an assigned priority of 0. If, on the contrary, the message already exists, meaning that the same sensor has already triggered one or more messages, then the priority of this existing message is incremented. During the decision stage, the message with highest priority on the list will be removed and processed against a table of rules to generate actions.

The table of rules. It describes a set of actions and their indices. The rules are initially similar for all characters. Each rule contains three parameters: index, action and priority. The index is again composed from the grammar $\{0, 1, \#\}$ and is used to match rules to a corresponding selected message being processed. Multiple rules can match one single message. The character #, which functions as a wildcard, implies that any value can be accepted for the corresponding particular character of the index. Furthermore, each of the rules has an assigned priority (initialized with a random value), and thus from all the candidate rules (with indices matching a selected message) only the one with highest priority is selected. The action to perform is coded on the second section of the rule, an alphanumeric code to be translated into a procedural action, such as an instruction to prey on any character that is within a certain distance, or an instruction to move towards the closest character, and so on (Figure 3).

The reward system. The priority of the rules is updated according to the consequences of the actions performed in the world. This translates into the character recognition of which actions are advantageous. An *ad hoc* reward was attributed to some of the possible actions such as eating when hungry, or victory or defeat on battle. If for instance the selected action is to feed, this implies a positive reward. On the contrary, being hit implies a negative value. The reward can be associated not only with the current rule which has triggered an event, but also preceding rules leading to the current action. Each character has a FIFO (First In First Out) memory stack which stores the last five rules performed. This block of memory is also rewarded accordingly, with the rules being

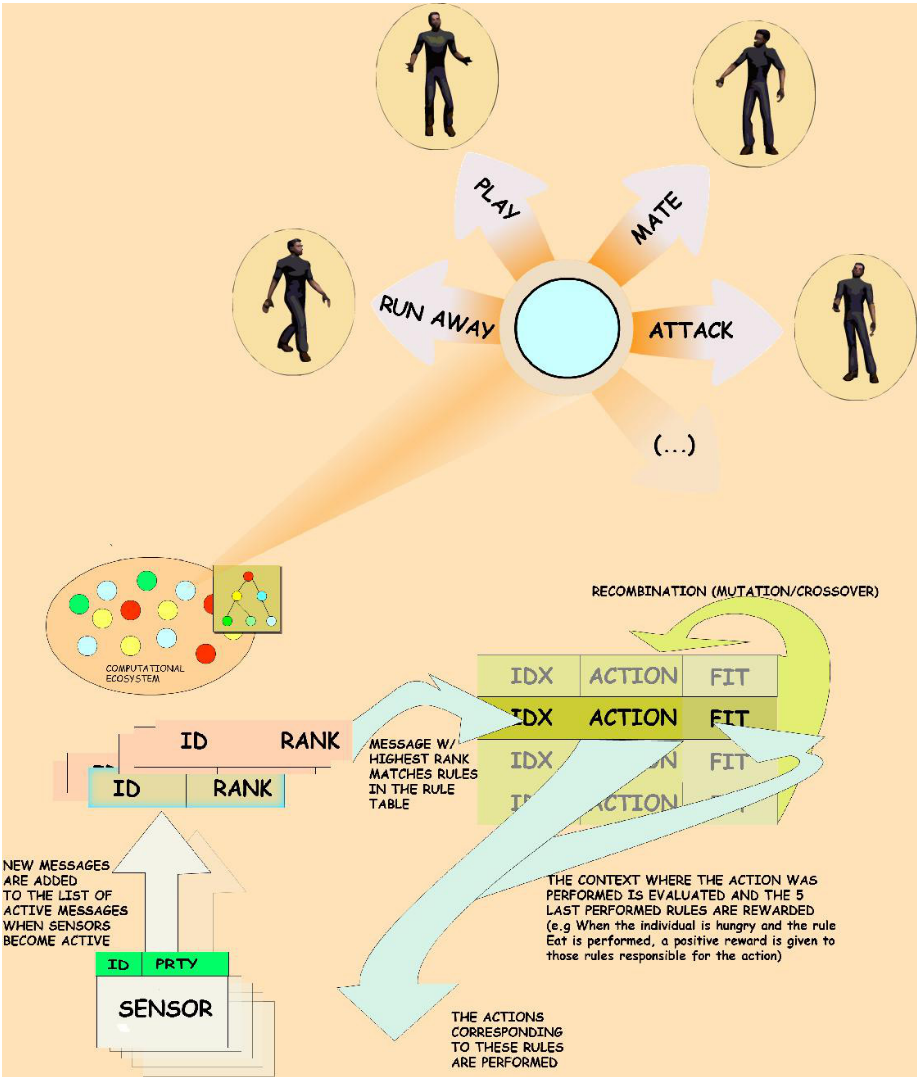

Fig. 3. Graphical depiction of the overall structure of the algorithm. On the bottom the action selection mechanism is illustrated: sensors trigger messages indicating a specific need; these messages are prioritized and form a buffer, which is ordered by associated priority. The message with the highest priority is selected to trigger an action. On the top of the graph the mechanism rendering actions is illustrated. For each syntactic action, an associated animation is played.

credited in a decremental way corresponding to the time they have been present on the memory. For instance, when a character manages to prey, the rule for the action which has triggered the preying event is rewarded with 5. The immediate rule–action prior to that one is rewarded with 4; the anterior with 3, and so on. When a new event occurs and a new action is performed, the oldest rule-action is removed from the stack.

Generation of new rules. As an outcome of reproduction, a newborn inherits from the current rule-table of the parents. To constitute the new rule-table, the rules with top priority from both progenitors are inherited in a 50/50 proportion. Each of the indices of the new set of rules then is perturbed by a process of (possible) mutations. These indices may suffer a transformation resulting from four possible attempts for digit mutation, each with a success probability of 50%. We set the mutation rate at such a high probability to ensure rapid variability in behaviors.

The Mapping of Behaviors. As mentioned earlier, actions are encoded in the second part of each rule which might trigger new messages or generate some physical action such as move or talk. To render visible each physical action in the virtual world one associated animation needs to be played. The relationship between the animations and the actions is rigidly defined *a priori*. For instance, for the rule triggering the action 'eat', the animation associated with 'eating' is performed (Figure 4 a)). However, these animations are not literal visualizations of these internal states, but rather were selected as interesting and playful in illustrating moments of conversation via gesticulating using well defined separate sets of body movements. The continuous dynamics of the virtual world is generated by the on-going displacements of the characters. These movements can be of three types: (i) in the direction of an arbitrary 'preferred location', *i.e.* a randomly selected layout point set of coordinates which, once reached, is re-initialised with a new target position; (ii) in the direction of other characters as a consequence of the internal needs as determined by the classifier system under the influence of the hormonal and metabolic sub–systems; and (iii) moving towards a member of the public in response to being selected for interaction (*i.e.* telling a story).

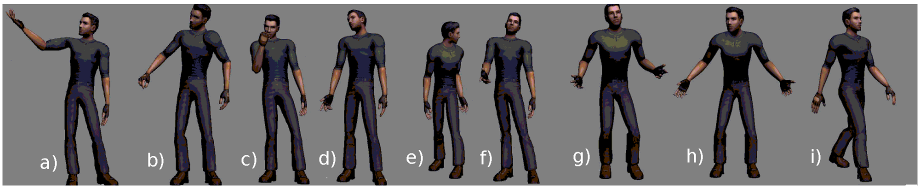

Fig. 4. Still images of the animations corresponding to each of the actions performed by the agent. From left to right: a) Eat a prey; b) attempt to mate with a partner that is not ready; c) reproduction or play alone; d) successful mating; e) attempt to mate but no compatible individuals exist on the vicinity; f) loosing an attack; g) play with happy mate; h) victorious attack; i) walking (move to mate, wander, move to prey).

3 Discussion

The *internal dynamics* of a functioning ecosystem is proposed as a way to structure and coordinate the animation of a population of humanoid characters. This *visualization* is exemplified with a population of NPCs simulating conversational behavior in the virtual world WisLM. The traditional approach of evolutionary art is to directly visualize the information defined on the Gtypes. In our work the information carried on the Gtype

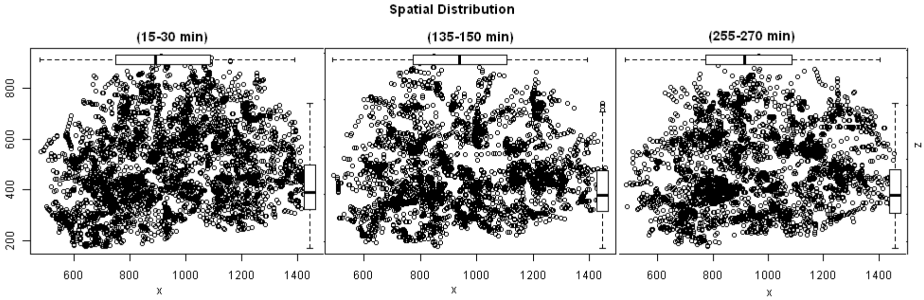

Fig. 5. Spatial distribution in time. At intervals of 15 seconds a snapshot of the system was captured, during a period of 6 hours (one typical day of exhibition). Each dot represents the location at that particular moment for one individual. The variables x and z stand for the coordinates in the Cartesian horizontal space. Each graph portrays a frame of juxtaposed 15 minutes of execution.

of the individuals describes features such as their body-composition and dietary constraints. However, in contrast to other approaches, what is emphasized and visualized in our system are the ephemeral states of the individuals within an ecosystem: *i.e.* the continuously changing actions they perform during their life-time reflecting the inner states of the CE, such as exchanges of energy. Similarly to the Gtype–Ptype paradigm from evolutionary computation, during the process of translation the linearity and the distance between the syntactic and the semantic levels can vary. For each action there is a corresponding symbol in the internal system and there is a posterior interpretation and visualization of that particular symbol. During this process of interpretation, converting symbols into visualizations, there is scope for creativity. The present research explores the ecosystem paradigm and the generative features implicit in CEs as a model of AI to develop new approaches for the animation of NPCs.

The result of this exploration is a populated landscape where individuals roam through the city, in a mapping of movements which is not random, with distinct emerging spatial attractor loci (Figure 5). Moreover there is place for heterogeneity and spontaneity and while ignoring some of the members of the crowd these autonomous individuals get together and aggregate in small groups in apparent expressive gesticulating dialogues (Figure 6). This system was exhibited to a public audience at the closing ceremony of the Watermans Festival of Digital Art 2012, and at the Tin Shed space, both in London, UK (Figure 7), with positive feedback from the audience. We obtained twenty four responses from anonymous users amongst the audience attending the shows. From an analysis of the responses to a questionnaire, 80% said the group formation and the talking movements in the simulation appeared to obey to an internal and coherent logic, whereas 65% said the behaviors to be believable in a cartoonish way.

To build our model of AI we have put together a set of established and efficient techniques for the agency of populations: Jones' hormone-framework, Saruwatu's model with dietary specifications and Holland's classifier system. Eventually, the complexity of the system could be reduced, *e.g.* events could be regulated by energy levels only. Our current design of a CE generates more noise in the system and consequent variability, than simpler implementations. Also, it provides a CE platform which is closer to a biology-based model, which we propose is of greater interest for an artwork such as WisLM which has its sources in human socio-history (of recent European colonialism). Nevertheless, our approach to a CE implementation is currently not computationally tractable for large crowds, when relying on the limited computer power of a single PC (our case and a typical situation when performing exhibits at small venues) — our model implemented a population restricted to a maximum of 200 simultaneous characters while still achieving real-time responses.

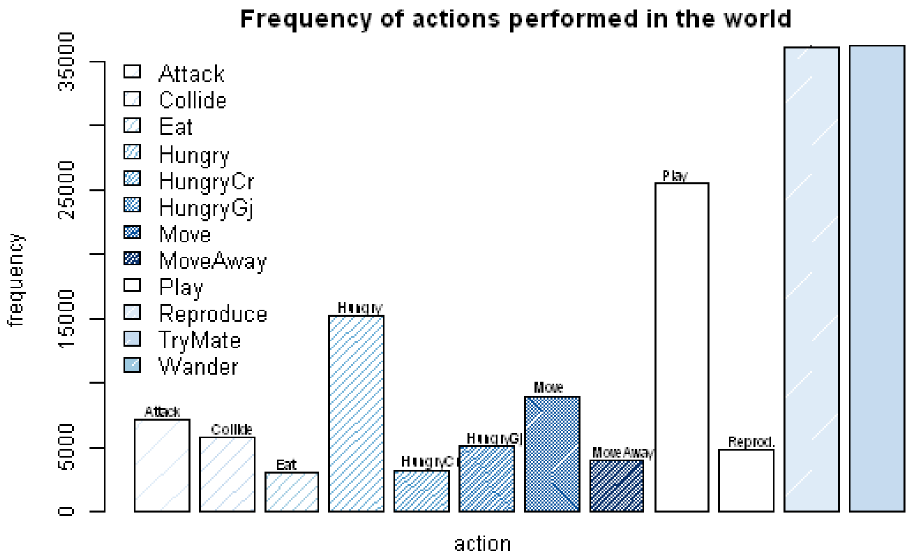

Fig. 6. The graph illustrates the number of times each action types was performed. At intervals of 15 seconds the system recorded what an individual was doing. From left to right, the actions are: Attack, Collide, Eat, Hungry, HungryCr (search for Carcasses), HungryGj (search for other creatures), Move, MoveAway, Play, Reproduce, TryMate, Wander.

Another aspect to take notice of, is that there is an implied semantic resulting from interactions of physical gestures which in our case was ignored. The resulting conversation is not immediately intelligible from the point of view of the actions in the ecosystem. Also, the small-number of pre-selected movements, as well as the lack of sophisticated blending which is limiting, hide the richness that could be effectively reached. An increased set of animations would make the system more flexible. To explore this potential further, it would be interesting to consider gestures and poses which better reflect the richness and combinatorics of the internal states of the characters. This could be enriched with a wider set of animations, which might reflect such nuances. Moreover, in contrast with our current deterministic approach of defining a limited set of animations, it would be interesting to break the animations into smaller bits. To explore the CE further, in terms of emergence and novelty this work would become richer with the incorporation of elements of a language of movements on which the CE could

(a) *(b)*

Fig. 7. *(a)* A perspective on the installation of WisLM during the exhibition at the Tin Shed, in London, in 2012. *(b)* A member of the audience interacts with one of the characters.

act. This would create the conditions for procedural movements with potential emergence of unexpected movements. Another interesting possibility would be of exploring the sonification of the CE internal states.

One advantage of the ecosystem approach to the animation of NPCs, over other established methods of population simulation, is its generative capacity. The system allows for a flexible representation of complex behavioral scenarios, and is relatively easily adaptable to different environmental contexts (*e.g.* buildings of various layouts and complexity levels) and varying numbers of individuals. The term "ecosystem" itself, and the associated terminology, are operative metaphors. The nature of the system is quite malleable. For example, instead of instances such as 'energy' the model admits variants such as 'currency', and then, events such as 'attack' or 'eat' could be mapped to become 'negotiate' or 'acquire'. Additionally, CEs, drawing on the ecosystem paradigm, provide natural fluctuations in the population density which might prove interesting from the point of view of the realism of the simulation. The drawback of the framework is however the difficulty in precisely controlling the behaviors observed, as the CE is, by definition, a complex system with a few variables upstream influencing potentially very complex detailed behaviors downstream.

4 Conclusions

Traditional genetic algorithms and evolutionary art are characterized by a process of interpretation of symbols between the Gtype and Ptype (*i.e.* genotype and phenotype). The approach we describe visualizes instead the ephemeral states of the individuals within the dynamics of an ecosystem implementation. The behaviors of the individuals during their normal activity, their primal movements and actions such as attack, flee, mate, prey, are taken as syntactic elements during a process of re-interpretation. For instance, the action of 'eating' produced at the inner ecosystem level might be translated, at the semantic level, as the animation of a wild animal chewing a carcass in the virtual world. However, as it happens with the original Gtype–Ptype paradigm, this process is also open to creativity and the linearity and distance of translation are subject to interpretation. For example, the action 'attack' rather than the animation of an animal fighting might instead correspond to a specific movement being choreographed to be performed by a dancer or, as determined in WisLM, might correspond to a conversational movement.

Our work explores these ideas with the animation of a population of humanoids NPCs in a virtual world playing the role of storytellers. However, the same specifications did also require the behavior of this community of humanoids to be the result of the dynamics of an ecosystem where individuals would seek out each other with predatory intentions. A CE, a system of agents organized in a hierarchical structure (of a food-chain) and trading units (of energy and biomass) generates such dynamics. The potential complexity of crowd interaction is revealed by the spontaneity of group formation of characters engaging in heterogeneous gesticulations during their conversations. The resulting populations offer spatial and behavioral distributions which are realistically far from uniform (Figures 5 and 6).

We have modeled a population of gregarious NPCs showing some of these spontaneous and conversational behaviors. Our approach took advantage of one of the fundamental properties of CEs: by relying on the variety and spontaneity of the elementary behaviors, the autonomy and self-organization of the agents generates ever-changing heterogeneous patterns at the global scale of a community. In other words, to build up this work we drew on the fact that the CE is, in essence, a dynamic generative framework which we have shown can be applied to animate NPCs.

Acknowledgments. This research is supported in part in the form of a PhD studentship to Mr. Antunes by the *Fundação para a Ciência e Tecnologia* from Portugal, contract reference SFRH / BD / 61293 / 2009. We also thank the telecom company Telefonica and in particular the VIDA competition for financial assistance under the award *Incentives for Ibero-American Production*. Finally, we thank the anonymous reviewers for their helpful comments and constructive criticism. Images 7a and 7b appear courtesy of Tin Shed Gallery.

References

1. Dorin, A.: Pandemic — Generative Software Installation, Exhibited: Bouillants 4, Vern-sur-Seiche, Brittany, France, Allin & Dupuis (artistic directors) (2012)
2. Eldridge, A., Dorin, A.: Filterscape: Energy Recycling in a Creative Ecosystem. In: Giacobini, M., et al. (eds.) EvoWorkshops 2009. LNCS, vol. 5484, pp. 508–517. Springer, Heidelberg (2009)
3. Antunes, R.F., Leymarie, F.F.: Generative Choreography: Animating in Real-Time Dancing Avatars. In: Machado, P., Romero, J., Carballal, A. (eds.) EvoMUSART 2012. LNCS, vol. 7247, pp. 1–10. Springer, Heidelberg (2012)
4. Shao, W., Terzopoulos, D.: Populating Reconstructed Archeological Sites with Autonomous Virtual Humans. In: Gratch, J., Young, M., Aylett, R.S., Ballin, D., Olivier, P. (eds.) IVA 2006. LNCS (LNAI), vol. 4133, pp. 420–433. Springer, Heidelberg (2006)
5. Maïm, J., Haegler, S., Yersin, B., Müller, P., Thalmann, D., Van Gool, L.J.: Populating Ancient Pompeii with Crowds of Virtual Romans. In: VAST, pp. 109–116 (2007)
6. Bogdanovych, A., Ijaz, K., Simoff, S.: The City of Uruk: Teaching Ancient History in a Virtual World. In: Nakano, Y., Neff, M., Paiva, A., Walker, M. (eds.) IVA 2012. LNCS, vol. 7502, pp. 28–35. Springer, Heidelberg (2012)

7. Huerre, S.: Agent–Based Crowd Simulation Tool for Theme Park Environments. In: 23rd Inter. Conf. on Comput. Anim. & Social Agents (CASA), Bournemouth University (2010)
8. Helbing, D.: A Fluid Dynamic Model for the Movement of Pedestrians. Complex Systems 6, 391–415 (1992)
9. Banerjee, B., et al.: Advancing the layered approach to agent–based crowd simulation. In: 22nd Workshop on Principles of Adv. & Distrib. Simulation, PADS, pp. 185–192 (2008)
10. Reynolds, C.W.: Flocks, Herds and Schools: A Distributed Behavioral Model. ACM SIGGRAPH, Computer Graphics 21(4), 25–34 (1987)
11. Pelechano, N.: et al.: Being a part of the crowd: Towards validating VR crowds using presence. In: Autonomous Agents & Multiagent Systems (AAMAS), pp. 136–142 (2008)
12. Pelechano, N., et al.: Virtual Crowds: Methods, Simulation, and Control. Synthesis Lectures on Computer Graphics and Animation. Morgan and Claypool (2008)
13. Thalmann, D., Musse, S.R.: Crowd Simulation. Springer (2007)
14. Jones, D.: AtomSwarm: A Framework for Swarm Improvisation. In: Giacobini, M., et al. (eds.) EvoWorkshops 2008. LNCS, vol. 4974, pp. 423–432. Springer, Heidelberg (2008)
15. Saruwatari, T., Toqunaga, Y., Hoshino, T.: ADIVERSITY: Stepping up trophic levels. In: 4th International Workshop on the Synthesis and Simulation of Living Systems, pp. 424–429 (1994)
16. Holland, J.: Hidden Order: How Adaptation Builds Complexity. Helix Books (1996)
17. McCormack, J.: Eden: An Evolutionary Sonic Ecosystem. In: Kelemen, J., Sosík, P. (eds.) ECAL 2001. LNCS (LNAI), vol. 2159, pp. 133–142. Springer, Heidelberg (2001)
18. Forrest, S., Jones, T.: Modeling Complex Adaptive Systems with Echo. In: Stonier, R., Yu, X. (eds.) Complex Systems: Mechanisms of Adaptation, pp. 3–21. IOS Press, Amsterdam (1994)