Entertainment Computing 5 (2014) 297–312

Contents lists available at ScienceDirect

Entertainment Computing

journal homepage: ees.elsevier.com/entcom

Semantic crowds ☆

Nick Kraayenbrink [1], Jassin Kessing, Tim Tutenel, Gerwin de Haan, Rafael Bidarra *

Computer Graphics and Visualization Group, Delft University of Technology, Delft, The Netherlands

ARTICLE INFO

Article history:
Received 30 May 2013
Revised 27 November 2013
Accepted 10 December 2013
Available online 18 January 2014

Keywords:
Crowd simulation
Crowd specification
Reusable crowds
Agent behavior
Semantic virtual worlds

ABSTRACT

Recent advances in crowd simulation techniques have led to increasingly realistic agent and group behavior. As many crowd simulation solutions typically target only specific types of environments and scenarios, numerous special-purpose methods and systems have emerged that are unsuitable for other contexts. Solving this situation demands a higher-level approach that takes re-use and re-configuration of crowds as a priority, for adequate application in a broad variety of scenarios, virtual environments and interaction with the entities present in that environment. In this article, we propose *semantic crowds*, a novel approach that allows one to specify and re-use the same crowds for virtually any environment, and have them use the objects available in it in a meaningful manner. To have the agents autonomously interact within any virtual world, we avoid in them explicit object-related information. Instead, this knowledge is stored in the objects themselves, which can then be queried, according to an agent's needs. To facilitate creating such crowds, we developed an interactive crowd editor that provides high-level editing parameters for defining crowd templates. We illustrate the flexibility of semantic crowds by means of three cases, in which we let the same crowd populate quite differently configured airport terminal environments. These examples also highlight that this modular approach easily combines with your custom implementations of agent behavior model and/or motion planner.

1. Introduction

Real-time crowd simulation is becoming increasingly important for a large variety of applications, such as games and simulation systems for training and education. For example, entertainment and serious games can expect a fair amount of criticism, if they feature rather unpopulated and empty cities and other environments. Recent advances in crowd simulation techniques have enabled more and more realistic agent and group behavior deploying, among other things, elaborate behavioral models, complex motion planning algorithms and impressive physics systems.

However, these improvements come at a price: most crowd simulation solutions are either too generic and high-level, or they are pretty much 'hard-wired' and customized, typically targeting only a few specific types of environments and scenarios. The former likely leads to repetitive, 'canned behavior', the latter to an ad hoc, specialized crowd, only suited for a very narrow application domain. In both cases, we are faced with a large variety of special-purpose methods and systems that are hard to re-configure and re-use in other contexts.

Avoiding these drawbacks requires a totally different approach, one that primarily supports and promotes the specification, configuration and re-use of crowds. In this article we propose one such approach, called *semantic crowds*, that has two main characteristics: (i) it allows one to specify once, and re-use with minimal modifications, the same crowds in virtually any environment; and (ii) in each environment, the agents in these crowds are able to use whatever objects are available in a meaningful manner.

We argue that, in order to successfully apply such specifications, or *crowd profiles*, in a broad variety of scenarios and virtual environments, the latter have to contain more than just pure geometry: they have to be extended with *semantics*. In the fields of linguistics, computer science and psychology, semantics is the study of meaning in communication. When focusing on virtual environments, we call semantics 'all information conveying the meaning of a virtual world and its entities' [1]. In this definition, entities encompass everything that can exist inside a world. Semantic virtual worlds, thus, consist of entities that 'know' not only about their shape and materials, but also about their attributes, roles, functions, services, etc. Such entities have the potential to strongly improve the quality and variety of interactions with an avatar, be it of a player or of any agent from a crowd [2].

In order to have the agents in a semantic crowd plausibly and autonomously interact within any virtual world, we minimize the information in the agents relative to what objects do and how to use them. Instead, that information is stored in the objects

☆ This paper has been recommended for acceptance by Matthias Rauterberg.

* Corresponding author. Tel.: +31 15 2784564.
E-mail address: r.bidarra@tudelft.nl (R. Bidarra).
[1] Current affiliation: Google Inc., Mountain View, CA, USA.

1875-9521/$ - see front matter © 2014 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.entcom.2013.12.002

themselves, which the agents can then purposefully query, based on what they want to achieve.

For the specification of crowds and agent behavior, we integrate several concepts, results and methods known from the literature. The novel contribution of our work lies mainly in that the same crowd, once specified, can be applied to whatever different virtual environments, while its agents always exhibit plausible yet very different behaviors without having to be 're-wired' themselves. In other words, we are making such crowds effectively reusable.

This article is structured as follows. We first survey previous work related to the specification of both virtual environments and crowds (Section 2). Next we elaborate on the essential aspects of semantics, particularly on its contribution to crowd and environment definition (Section 3). We then introduce the main concepts of the semantic crowd model, both in its demographics and in its agents (Section 4), and describe the main features of our prototype framework, including an interactive crowd editor that provides high-level editing parameters for defining crowd profiles (Section 5). Finally, we illustrate the potential and flexibility of semantic crowds by means of three demonstration cases (Section 6).

2. Related work

This section surveys a selection of previous work regarding the representation and generation of virtual environments and crowds.

2.1. Virtual environments

Beyond geometric data, virtual environments may also contain information about objects, which actions these can perform or how they can be used by agents. A Virtual Environment (VE) presenting this kind of information is normally referenced as an Intelligent Virtual Environment [3], Informed Environment [4] or Semantic Virtual Environment [5].

Thomas and Donikian [6] propose a model of virtual environments using structures suitable for behavioral animations. Using this knowledge, autonomous virtual actors can behave like pedestrians or car drivers in an urban environment.

Through the concept of synoptic objects [7], Badawi and Donikian describe an informed environment using objects which contain a synopsis of interactions that they can be subjected to. Using a set of seven basic actions, the objects can describe the interaction process to any agent using them. Smart objects [8] were a successful proposal for adding semantics to virtual objects, dealing with many of the possible user interactions in a VE. They were primarily devised for manipulation, animation, and planning purposes. Gutierrez, Vexo and Thalmann [9] present an object representation based on the semantics and functionality of interactive digital items within a VE. Each object participating in a VE application is not only a 3D shape, but a dynamic entity with multiple visual representations and functionalities. In this way, it is possible to dynamically scale and adapt the object's geometry and functions to different scenarios.

Research in artificial intelligence proposed the notion of ontologies, to overcome the lack of shareable and reusable knowledge bases [10]. Ontologies define the meaning of objects and the relations between them. Using this concept to define a world knowledge base, Grimaldo et al. [11] propose a semantics-based framework for simulation of groups of intelligent agents. Agents can use the information provided to enhance both agent-object and agent-agent interactions. An object taxonomy is used to classify the interactive objects and to specify their properties. Ontologies are also used to define social relations among agents in order to display socially acceptable decisions.

A semantic model for representing multi-layered complex environments is presented by Jiang et al. [12] and is composed of three different levels: a geometric level, a semantic level, and an application level. The geometric level contains a 3D model of the environment that is used for visualization and to extract semantic information that will feed the next semantic level. This semantic level comprehends a structure map, a topologic map and a height map which are used to identify or query semantic information of the environment. The final layer (application level) is responsible for providing efficient interaction between pedestrians and the environment. The crowd model used in their work is a modified version of the original model proposed by Treuille et al. [13].

2.2. Virtual crowds

Virtual groups have been studied since the early days of behavioral animation. There are two main approaches used in the literature: agent-based models and force-based models. The first approach is based on modeling of virtual agents, which interact among themselves and have some level of autonomy and individuality. The second approach -force-based models- provides more global control and handles high density crowds. These methods normally yield crowd simulations that resemble particles rather than human animation. In this subsection, we survey existing approaches for crowd models.

Two seminal papers are examples of agents-based model: Reynolds [14] simulated flocks of bird-like entities, or boids, obtaining realistic animations by using only simple local rules; Tu and Terzopoulos ([15]) created groups of artificial fishes endowed with synthetic vision and environmental perception, which control their behavior. In both papers, only small groups were simulated, and high density crowds were not treated. Another example of agents-based models was presented by Musse and Thalmann [16], who proposed an approach with hierarchically structured crowds having different levels of autonomy. In their model, the behavior is based on a set of rules dealing with the information contained in groups of individuals, e.g. individual and group knowledge about the world, individual and group states as well as their intentions. Ulicny and Thalmann [17] proposed a model for crowd simulation based on combination of rules and Finite State Machines for controlling agents' behaviors in a multi-layer approach. The model proposed by Farenc et al. [18] describes the coordination between smart objects, intelligent environments and the virtual crowd. This work describes crowds of virtual people which were controlled by the objects in order to act, and could read the environment in order to evolve in the simulation. In this case, the virtual agents were less autonomous and more controlled by the environment. Not focused on crowds, but on groups of a few virtual agents, the model proposed by Abaci et al. [19] describe an extended version of smart objects for AI and planning purposes.

More recently, Kapadia et al. [20] proposed a multi-actor simulation in which actions, modifiers, constraints and behaviors can be expressed in a scripting language. These are used in a behavior state machine for the actors. To actions, cost effects can be defined and by expressing a minimize or maximize on particular costs in a behavior the most desired actions are chosen. Allbeck [21] introduced the CAROSA tool for authoring NPCs. The framework contains an 'Actionary' which includes a database of parameterized actions and objects. Both can be defined independent of a particular application or scenario. These actions and objects are built up hierarchically. A scheduler can be used to link particular actions to a time and place. Stocker et al. [22] introduce an approach where dynamically occuring events describe what types of actions their virtual humans can execute. Additionally they introduced the idea of agent priming to further restrict and simplify action choice.